

HP

Storage Essentials 5.0 User Guide



T3710-96002

Part number: T3710-96002
First edition: September 2005



Legal and notice information

© Copyright 2005 ApplQ Inc.

© Copyright 2005 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company and ApplQ Inc. makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard and ApplQ shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information is provided "as is" without warranty of any kind and is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP and ApplQ shall not be liable for technical or editorial errors or omissions contained herein.

Windows are registered trademarks of Microsoft Corporation. Oracle is a registered trademark of Oracle Corporation. Sun, Solaris, Sun StorEdge, and Java are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. AIX and IBM are registered trademarks of International Business Machines Corporation in the United States, other countries or both. SGI and IRIX are registered trademarks of Silicon Graphics, Inc. Netscape is a registered trademark of Netscape Communications Corporation in the United States and other countries. HDS and HiCommand are registered trademarks of Hitachi Data Systems. HP-UX is a registered trademark of Hewlett-Packard Company.

UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Other product and company names mentioned herein may be the trademarks of their respective owners.

Storage Essentials 5.0 User Guide

Contents

About this guide	xix
Intended audience	xix
Prerequisites	xix
Related documentation	xix
Document conventions and symbols	xx
HP technical support	xxi
HP-authorized reseller	xxi
Helpful web sites	xxi
1 Overview	1
About This Product	1
Suggested Topics for First-Time Users	4
Product Components	4
Management Server Components	5
About the User Interface	6
About the Top Pane	6
Installing the Java Plug-in	8
Installing the Software Security Certificate	9
Installing the Certificate by Using Microsoft	
Explorer 6.0	9
Installing the Certificate by Using Netscape	
Navigator 7	10
Changing the Security Certificate to Match	
the Name of the Server	10
Stopping the Service for the	
Management Server	11
2 Discovering Filers, Tape Libraries, Switches, and Storage Systems.	13
Step 1 - Discover Switches	14
Discovering Brocade Switches.	15
Verifying Brocade Rapid Program Is Set to 1	15
Discovering CNT Switches	16
Discovering Cisco Switches.	17
Discovering Sun StorEdge and QLogic Switches	17
Changing the SNMP Trap Listener Port for Sun StorEdge Switches	18
Discovering McDATA and EMC Connectrix Switches	18
SWAPI Setting Through a Proxy	20
Step 1 - (McDATA Switches Only) Install the Bridge Agent	20
Step 2 - Change the Discovery Setting for McDATA and Connectrix Switches to SWAPI	21
Step 3 - Discover the Proxy	21
CIM_ERR_FAILED When Trying to Activate a Zone Set Using McDATA SWAPI	21

SNMP Setting Through a Proxy.	22
Step 1 - Verify the Discovery Setting for Switches Is Set to SNMP.	22
Step 2 - Discover the Proxy	22
Step 3 - Make Sure There Are No Port Conflicts for Receiving SNMP Traps.	23
Step 4 - Step Up the Proxy to Send Traps to the Correct Port	23
Contacting a McDATA or Connectrix Switch Directly.	23
Make Sure There Are No Port Conflicts for Receiving SNMP Traps.	23
Configure the SNMP Agent to Send Traps to the Correct Port	24
Changing the Discovery Settings.	24
Excluding McDATA and EMC Connectrix Switches from Discovery	25
Viewing Log Messages.	26
Duplicate Logs for Brocade Switches in Same Fabric.	26
Step 2 - Discover Storage Systems, Filers and Tape Libraries	26
Discovering EMC Solutions Enabler 5.1	29
Excluding EMC Symmetrix Storage Systems from Discovery	30
Discovering EMC CLARiiON Storage Systems.	31
Discovering HDS Storage Systems	31
Excluding HDS Storage Systems from Discovery	32
Discovering HP StorageWorks Arrays	33
Discovering Engenio Storage Systems	34
Discovering NetApp Filers	34
Discovering Sun StorEdge 3510 Storage Systems	35
Discovering Sun StorEdge 6920 Storage Systems	36
Discovering Sun StorEdge 6130 Storage Systems	36
Discovering IBM Storage Systems	36
Discovering IBM Tape Libraries	37
Modifying the Properties of a Discovered Address.	37
Deleting Elements from the Management Server	38
Deleting an Element Using System Manager or Chargeback Manager	39
Step 3 - Discovery Data Collection	39
Discovery Data Collection.	39
Stopping the Gathering of Details	41
Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh	41
Excluding HDS Storage Systems from Force Device Manager Refresh.	43
Managing McDATA and EMC Connectrix Switches	44
About Managing McDATA and EMC Connectrix Switches	44
Adding McDATA and EMC Connectrix Switches	44
Assigning a File Extension in Netscape 7	44
Updating the Database with Element Changes	45
Filtering Discovery Groups.	45
Moving Elements to Another Discovery Group	46
Placing an Element in Quarantine.	46
Removing an Element from Quarantine	47
3 Discovering Applications and Hosts	49
Step 1 - Discovering Your Hosts	49

Step A - Set Up Discovery for Hosts	50
Enabling Dynamic Disk Detection for Windows 2000 Hosts.	50
Step B - Obtain Details	50
Step 2 - Setting Up Discovery for Applications	51
Monitoring Oracle	52
Step A - Create the APPIQ_USER Account for Oracle	52
Removing the APPIQ_USER Account for Oracle	53
Step B - Provide the TNS Listener Port	54
Step C - Set up Discovery for Oracle 10g	55
Discovering Oracle Clusters	56
Monitoring Microsoft SQL Server	57
Switching to Mixed Mode Authentication.	57
Step A - Create the APPIQ_USER for the SQL Server	57
Removing the APPIQ_USER Account for SQL Server	58
Step B - Provide the Microsoft SQL Server Name and Port Number	59
Deleting SQL Server Information	59
Monitoring Sybase Adaptive Server Enterprise	60
Step A - Create the APPIQ_USER account for Sybase	60
Removing the APPIQ_USER Account for Sybase	61
Step B - Provide the Sybase Server Name and Port Number	62
Deleting Sybase Information.	63
Monitoring Microsoft Exchange	63
Adding Microsoft Exchange Domain Controller Access	63
Deleting a Microsoft Exchange Domain Controller	64
Step 3 - Discovering Applications	64
Step A - Detect the Applications.	65
Step B - Obtain Details	65
Changing the Oracle TNS Listener Port	66
Adding/Modifying Microsoft Exchange Domain Controller Access	66
Deleting a Microsoft Exchange Domain Controller	67
Changing the Password for the Managed Database Account	67
Obtaining Disk Drive Statistics from Engenio Storage Systems	68
Assigning a File Extension in Netscape 7	69
4 Managing Security	71
About the Security for the Management Server	71
About Roles	71
About Organizations	74
Planning Your Hierarchy.	77
Naming Organizations	77
Managing User Accounts	78
Adding Users	78
Editing a User Account.	79
Deleting Users.	80
Modifying Your User Profile	80
Modifying Your User Preferences.	81
System Manager and Element Topology Preferences	81

Event Monitoring for Storage Essentials Preferences	81
Warnings for Slow Systems Operations	82
Viewing the Properties of a Role	82
Viewing the Properties of an Organization	83
Managing Roles	83
Adding Roles	83
Editing Roles	84
Deleting Roles	85
Managing Organizations	86
Adding an Organization	86
Viewing Organizations	87
Editing Organizations	88
Deleting an Organization	88
Removing Members from an Organization	89
Accessing the Edit Organization Window	89
Accessing the Add Elements to Organization Window	90
Filtering Organizations	90
5 Managing Licenses	93
About the License	93
Importing a License File	96
Viewing Cumulative Licenses	96
Viewing a Specific License	97
Deleting a License	98
6 Configuring the Management Server	99
Trap Generation	99
Setting Up E-mail Notification	99
Managing Getting Discovery Details	100
Adding a Discovery Schedule	100
Disabling a Schedule	102
Editing a Schedule	102
Removing a Schedule	102
Managing Product Health	102
Enabling Disk Space Monitoring	103
Viewing the Results of Disk Space Monitoring	104
Advanced Settings	104
Modifying Java Memory Settings	105
Customizing Properties	105
Managing Logging	105
Accessing the Log Files	106
Downloading Logs to a File	106
Downloading the User Audit Log	107
Downloading the Discovery Summary Log	107
Displaying a Log File in a Command Prompt Window	107
Changing the Provider Log Level	108
Enabling the Scanning of Critical Events	

of the Management Server Database	109
Viewing the Results of Logging	110
Managing the Display of Events	110
Controlling the Display of Cleared and Deleted Events	110
Modifying the Clearing and Deletion Frequency	111
Configuring the Clearing of Events.	112
Configuring the Deletion of Events	113
Managing File System Viewer	113
Managing Collectors for File System Viewer	114
Setting Data Collection Properties	115
Scheduling File System Viewer Scan Checker and Scan Status	116
Performing a Test File System Viewer Scan	118
Managing Reports	118
Architectural Overview of Report Views and Report Cache Refresh	118
Managing Collectors for Reports	121
Starting Collectors.	122
Stopping Collectors.	122
Setting the Date and Time for Report Collectors	123
Viewing Scheduled E-mail Deliveries for Reports	124
Editing E-mail Schedules for Reports.	125
Viewing Data Aging Statistics for Reports	126
Scheduling Report Cleanup.	127
Refreshing the Report Cache	128
Refreshing the Report Cache Immediately	128
Scheduling a Report Cache Refresh.	129
Setting Up Global Reporter.	130
Editing Remote Site Information	135
Remote Sites Are Not Removed from the tnnames.ora File	136
Managing Performance Collection	136
Managing Performance Collectors	137
Starting Performance Collectors	138
Stopping Performance Collectors	139
Setting the Date and Time for Performance Collectors.	139
Viewing Data Aging Statistics for Performance	140
7 Database Maintenance and Management	143
Database Maintenance Window	143
Overview of Backups	144
Database Mode	145
Architectural Overview of RMAN Backups.	146
Performing an RMAN Hot Backup	148
Scheduling RMAN Hot Backups of the Database	148
Viewing Results from RMAN Backup.	149
About the Database Admin Utility	150
Accessing the Database Admin Utility	150
Refreshing the Database Admin Utility	150
Checking the Database and Listener Status	151

Changing the Password of System Accounts	151
Exporting the Database	152
Importing the Database	153
Reinitializing the Database	153
Resetting the Temp and Undo Tablespace	154
Restarting the Database	154
Restoring a Cold Backup	154
Setting the Archive Mode	155
RMAN Restore	156
Running a Cold Backup	156
Changing the Archive Destination	157
Viewing Logs from the Database Admin Utility	157
Resetting the Log	158
Warning Messages During Reinitializing the Database	158
Database Password	158
Troubleshooting Listener and Database Connection Problems	159
8 Viewing Element Topology and Properties	161
About System Manager	161
Accessing System Manager	163
About Cisco Switches and VSANs in System Manager	163
About the User Interface	164
About the User Interface for System Manager	164
About the Toolbar in System Manager	164
Description of the Icons Displayed in the Topology	167
About the List Tab	167
Viewing Elements by Type	168
About the Access Tab	169
Obtaining Information About Zone Entries	170
Obtaining Information About Host Bindings	172
Obtaining Information About Storage System LUN Masking	173
About the Path Tab	175
About the Right-Click Menu Options	176
Viewing Storage Elements	182
Adding a Virtual Application	183
Adding Information for “Discovered” Hosts	184
Arranging Elements in the Topology	185
Closing Topology Windows	187
Using the Global View	187
Printing Elements in System Manager	187
Updating Element Data	189
Viewing Ports	190
Showing the Impact of an Element	190
Assigning a Business Cost to an Application	192
Expanding the Topology Pane	194
Filtering Fabrics	194
Viewing Event Status in the Topology	194

Custom Name for Switch Truncated in Topology	197
Managing Groups	197
About Groups	197
Grouping "Discovered" Hosts.	197
Ungrouping "Discovered" Hosts	198
Grouping "Discovered" Storage Systems	199
Ungrouping "Discovered" Storage Systems	200
Managing Fabrics	200
Changing the Fabric Name	200
Deleting Fabrics	200
Hiding and Showing Generic Hosts	201
About Hiding Generic Hosts.	201
Hiding Generic Hosts for a Switch	202
Expanding Generic Hosts for a Switch	202
Hiding Generic Hosts for All Switches	203
Expanding Generic Hosts for All Switches	203
Setting Up Custom Commands	203
About Custom Commands	204
Adding a Custom Command.	204
Editing a Custom Command	207
Deleting a Custom Command	207
Software Environment Variables for Scripting	207
Using the Remote Console	210
About the Remote Console.	211
Keeping the Remote Console Active	211
About the Buttons on the Remote Console	212
About the Drop-Down Menu Options.	213
Copying Text from the Remote Console	213
Using External Tools	213
About the External Tools Feature	214
Setting up External Tools.	214
About the Navigation Tab	215
Finding the Status of a Port on a Switch	218
Accessing the Navigation Tab.	219
Viewing Element Properties	219
About the Properties Tab	219
Accessing the Properties Tab.	220
Viewing Fabric Properties.	221
Assigning a Custom Name	221
Adding a Generic Tape Library	222
Viewing Element Topology.	223
About the Topology Tab	223
Accessing the Topology.	228
About the New Window Option	229
Printing the Topology	230
Creating a Virtual Application	232
About the Provisioning Tab	233

About the Events Tab	233
Asset Attributes of an Element	234
About the Collectors Tab	236
About the Monitoring Tab	236
About the Policies Tab	237
Determining If a Host Belongs to a File System	237
About the Data from CXFS File Systems	238
9 Provisioning Manager	239
About Provisioning Manager	239
Managing Zones	240
SAN Zoning Overview	240
Accessing Information About Zone Aliases	244
Creating Zone Aliases	245
Modifying a Zone Alias	246
Deleting Zone Aliases	246
Accessing Information About Zoning	247
Creating a Zone in a Fabric	247
Adding and Removing Zone Members	248
Deleting Zones	249
Accessing Information About Zone Sets	249
Creating a Zone Set	249
Modifying a Zone Set	251
Deleting Zone Sets	251
Copying a Zone Set	252
Activating a Zone Set	253
Zones and Zone Sets Are Sometimes Listed Twice	254
Changing the Amount of Information Collected from the Inactive Zone Database (Cisco Switches)	255
Managing Storage	256
Setting Up Storage Partitioning	256
Modifying the Cache Settings (Engenio Only)	260
Changing the Owner of a Volume (Engenio Only)	260
Managing Storage Pools	261
Creating a Storage Pool (Engenio and CLARiiON Only)	261
Accessing Information About Storage Pools	262
Deleting a Storage Pool (Engenio and CLARiiON Only)	263
Managing Volumes	263
Accessing Information About Volumes	263
Volumes on HDS Storage Systems	264
Creating a Storage Volume	265
Deleting a Storage Volume	268
Changing the Cache Block Size for a Storage System	269
Modifying the Cache Settings (Engenio Storage Systems Only)	269
Rules for Creating Host Security Groups	270
Managing Host Security Groups	275

Accessing Information About Host Security Groups	275
Creating Host Security Groups	276
Editing Host Security Groups	279
Deleting Host Security Groups	281
Setting the Host Mode for IBM Storage Systems	281
General Provisioning Issues.	282
Provisioning Can Make a Device Inaccessible	282
Provisioning Does Not Make an Operating System Aware of a Device	282
Provisioning Issues by Vendor	282
Issues Specific to CLARiiON Storage Systems.	283
Make the Management Server a Privileged User for CLARiiON	283
EMC Navisphere CLI Is Required	283
Issues Specific to EMC Symmetrix Storage Systems	283
About Provisioning on EMC Symmetrix Storage Systems	284
LUN Provisioning Is Not Supported with Solutions Enabler 4.3	284
"Process Has an Exclusive Lock" Message.	284
Some EMC Volumes, Their LUNs and LUN Maskings Are Hidden	285
Issues Specific to HDS Storage Systems	285
About Provisioning on HDS Storage Systems	285
Unable to Provision When HDS CruiseControl Is Enabled	286
Increasing the Wait Time for the Management Server.	286
Initiator Ports Cannot Be Used for Provisioning	287
The Software Is Unable to Map a Host to Volumes in Different Host Storage Domains of the Same Port on HDS 9900V	287
A Default LUN Number Is Used Instead of a User-Specified One	287
Issues Specific to Engenio Storage Systems	287
Creating and Deleting Storage Pools	287
Creating and Deleting Storage Volumes	287
Creating a LUN	288

10 Path Provisioning. 289

About Path Provisioning	289
How Path Provisioning Works	291
How to Use Path Provisioning	291
System Actions.	293
Volume Creation, LUN Security, and Zone Operation	293
Step 1 - Select Storage System	293
Step 2 - Select a Host	294
Step 3 - Select a Volume	295
Step 4 - Select a Host Security Group	296
Step 5 - Select a Zone	296
Creating a Zone	297
Meta Volume Creation	298
LUN Security.	299

Step 1 - Select Storage System	300
Step 2 - Select a Host	301
Adding a Host.	302
Step 3 - Select a Volume	303
Step 4 - Select a Host Security Group	303
Zone Operation	304
Step 1 - Select Storage System	304
Step 2 - Select a Host	305
Step 3 - Select a Zone	306
Creating a Zone	307
Volume Creation and LUN Security	308
Step 1 - Select Storage System	308
Step 2 - Select a Volume	309
Step 3 - Select a Host Security Group	310
Providing a LUN Number	310
Adding a Host	311
Creating a Host Security Group	311
Scheduling Provisioning Jobs	313
Executing Provisioning Jobs	314
Monitoring Provisioning Jobs	315
Deleting Multiple Jobs	316
Naming Conventions for Zones	316
Using Multipathing with Path Provisioning	317
Customizing Path Provisioning	317
Storage System Customize Dialog	318
Host Customize Dialog	318
Volume Customize Dialog	318
LUN Customize Dialog	318
Zone Customize Dialog	319
Automatically Configure Zoning	320
Manually Configure Zoning	320
11 Event Management	323
About Event Monitoring for Storage Essentials	323
Viewing Event Details	326
Clearing Events	327
Unclearing Events	327
Configuring the Clearing of Events	328
Configuring the Deletion of Events	328
Deleting Events	329
Sorting Events	329
Changing the CLARiiON Event Polling Interval	330
Brocade Events	331
Brocade Switch Events	331
Supported Brocade Events	332
Filtering Events	332
About Filtering Events	333

Selecting a Severity Level	333
Customizing the Severity Level Filter	334
Filtering Events by Summary Text	335
Hiding Filters in Event Monitoring for Storage Essentials	335
Enabling the Rank Column in Event Monitoring for Storage Essentials	335
12Running Reports	337
About Reporting	337
Accessing Reporting	339
Viewing Reports	340
Viewing Report Collectors for an Element	340
Refreshing a Report	341
Changing the Formatting of a Report	341
Opening a Report in a New Window	342
Maximizing the Screen Space for a Report	342
Filtering Data in Global Reports	342
Sending a Report by E-mail	343
Managing E-mail Schedules for Reports	344
Adding an E-mail Schedule for a Report	344
Editing an E-mail Schedule for a Report	347
Deleting E-mail Schedules for a Report	349
Viewing E-mail Schedules for a Report	349
Creating Custom Reports	350
About Creating Custom Reports	350
Configuring Report Designer to Work with the Management Server	351
Designing Custom Reports	353
Creating Standard Reports	353
Integrating Custom Reports	362
Detailed Schema Information	364
Views from Previous Releases	397
13Viewing Performance Data	399
About Performance Manager	399
Accessing Performance Manager	400
Creating Performance Charts	400
About the Toolbars in Performance Manager	401
Comparing the Performance of Different Elements	404
Viewing Summary Charts	404
Viewing Trending Information for Performance	405
Setting Thresholds in Performance Charts	406
Removing Performance Data from a Graph	406
Setting a Custom Period	407
About the Monitoring Options	408
Task Already Running on Element	417
Monitoring with Direct Attached Storage	418
Supported Host Configurations for Monitoring	418

Sudden Dips Displayed in Certain Charts in Performance Manager	420
Values Continue to Increase in Charts for Aggregated Drives and Aggregate Volumes	421
14 Finding an Element's Storage Capacity	423
About Capacity Manager	423
Accessing Capacity Manager	425
About the Toolbars in Capacity Manager	425
Finding the Capacity of an Element	428
Additional Information for Applications	429
Additional Information for Hosts and Filers	429
Additional Information for Storage Systems	430
Obtaining Utilization Reports	431
Printing Elements in Capacity Manager	431
Viewing Capacity Charts	433
Viewing Trending Information for Storage Capacity	435
Different Results for the df -k Command and Capacity Manager	435
15 Managing Policies	437
About Policy Manager	437
Accessing Policy Manager	438
Creating Policies	438
Creating a Utilization Policy	438
Creating Policies for Discovery	441
Creating Policies for Provisioning	442
Creating Policies for Events	444
Testing a Utilization Policy	446
Modifying Policies	446
Modifying Utilization Policies	446
Modifying Discovery Policies	448
Modifying Provisioning Policies	449
Modifying Policies for Events	451
Viewing Policies	453
Deactivating a Policy	453
Deleting Policies	453
Providing E-mail Notification for a Policy	454
Providing Event Generation for a Policy	454
Providing a Custom Command for a Policy	455
16 Chargeback Manager	457
About Chargeback Manager	457
Setting Up Chargeback Manager	458
Accessing Chargeback Manager	459
Creating an Asset Record	459
Changing the Status of an Element	460
Saving Chargeback Manager Information	461

Viewing Assets	461
Setting Storage Tier Name and Cost	462
Setting the Infrastructure Cost	463
Adding Asset Information	463
Adding Asset Information	463
Adding General Information	464
Adding Staff Information	465
Adding Geographic Information	466
Adding Licensing and Warranty Information	466
Adding Custom Information	466
Managing Departments	466
Adding Departments	467
Editing a Department	467
Removing a Department from Chargeback Manager	467
Setting Up Asset and Storage Based Chargeback Manager	468
Setting Up Asset-Based Chargeback Manager	468
Step 1 - Specify Financial information	469
Step 2 - Assign Departmental Ownership Percentage	470
Step 3 - Review Asset-based Chargeback Result	470
Setting Up Storage-Based Chargeback Manager	471
Step 1 - Assign Departmental Ownership Percentage	472
Step 2 - Review Storage Tier Cost	473
Step 3 - Review Storage Dependency and Cost	473
Step 4 - Review Storage-Based Chargeback Result	474
Editing Percentage of Ownership	475
Removing Department Ownership of an Element	475
Capacity in Chargeback Manager and Capacity Manager Differ	476
How a Depreciation Method Is Calculated	476
Calculating Straight Line Depreciation	477
Calculating Fixed Declining Balance	478
Calculating Double Declining Balance	479
Viewing Chargeback	481
Viewing Chargeback by Element	481
Viewing Chargeback by Department	482
Viewing Chargeback by Owner	483
Chargeback Reports	484
Viewing Chargeback Reports	484
E-mailing a Chargeback Report	485
Managing E-mail Schedules for Chargeback Reports	486
Adding an E-mail Schedule for a Chargeback Report	486
Editing an E-mail Schedule for a Chargeback Report	488
Deleting E-mail Schedules for a Chargeback Report	489
Viewing E-mail Schedules for a Chargeback Report	490
Viewing the History of an E-mail Chargeback Schedule	491
Filtering Assets	491
About Filtering Assets	492
Selecting an Element Type for Chargeback	492

Customizing the Element Type Filter for Chargeback	493
Filtering Assets by Status	493
Customizing the Asset Status Filter for Chargeback	493
Hiding Filters in Chargeback Manager	494
17 Business Tools	495
About the Business Tools	495
Using the HBA Replacement Automator	496
Installing New HBA with Old HBA	496
Installing New HBA by Itself	496
Using the Brocade Firmware Update Utility	497
Updating the Firmware Without Modifying the Configuration File	497
Updating the Firmware With Modifying the Configuration File	498
Setting up Risk Analysis	499
Global Change Management	499
18 Troubleshooting	501
"Data is late or an error occurred"	
Message	501
appiq.log Filled with Connection	
Exceptions	501
Receiving "HTTP ERROR: 503" When	
Accessing the Management Server	502
Errors in the Logs	503
Permanently Changing the Port a	
CIM Extension Uses (UNIX Only)	503
Configuring UNIX CIM Extensions	
to Run Behind Firewalls	504
Volume Names from Ambiguous	
Automounts Are Not Displayed	508
Installing the Software Security Certificate	509
Installing the Certificate by Using Microsoft Explorer 6.0	509
Installing the Certificate by Using	
Netscape Navigator 7	510
Changing the Security Certificate to Match the	
Name of the Server	510
Troubleshooting Discovery and	
Discovery Data Collection	511
Configuring E-mail Notification for Discovery Data Collection	511
Increasing the Time-out Period and Number of	
Retries for Switches	512
"Connection to the Database Server Failed" Error	514
DCOM Unable to Communicate with Computer	514
Duplicate Listings for Brocade Switches in Same Fabric	514
Element Logs Authentication Errors During Discovery	515
EMC Device Masking Database Does Not Appear	
in Topology (AIX Only)	515

Microsoft Exchange Drive Shown as a Local Drive	515
Unable to Discover Microsoft Exchange Servers	515
Nonexistent Oracle Instance Is Displayed	515
Requirements for Discovering Oracle	515
Unable to Find Elements on the Network	516
Unable to See Path Information	516
Device Locking Mechanism for Brocade Element Manager Query/Reconfiguration	516
A Discovered Sun StorEdge A5000 JBOD Does Not Display Its WWN Properly	516
Unable to Monitor McDATA Switches	516
Unable to Detect a Host Bus Adapter	517
Navigation Tab Displays Removed Drives as Disk Drives	517
Unable to Obtain Information from a CLARiiON Storage System	517
Discovery Fails Too Slowly for a Nonexistent IP Address.	517
"CIM_ERR_FAILED" Message	518
Communicating with HiCommand Device Manager Over SSL	519
Unable to Discover a UNIX Host Because of DNS or Routing Issues	520
Troubleshooting the Java Plug-in	521
Java Applet Has Data from a Different Version of Management Server Software	521
OutOfMemoryException Messages	522
Improving Reload Performance in System Manager	522
Troubleshooting Provisioning	522
Cannot Access a Resource Owned by Another Controller.	522
Error -56.	522
"Can't delete this zone" Message	523
Changes in EFC Manager Requiring "Discovery Data Collection"	523
Troubleshooting Hardware	523
About Swapping Host Bus Adapters.	523
"Fork Function Failed" Message on AIX Hosts	523
Known Driver Issues	523
Known Device Issues	524
"mailbox command 17 failure status FFF7" Message	527
"Process Has an Exclusive Lock" Message	527
Glossary	529
A	529
C	529
F.	529
D.	530
E.	530
G	530

H	530
I	530
M	531
P	531
S	532
V	532
W	532
Z	533
Index	535

Figures

1 Accessing Features from the Tools > Storage Essentials Menu	2
2 Accessing Roles in Storage Essentials	2
3 Accessing the Documentation Center	3
4 Status Light	7
5 Parent-Child Hierarchy for Organizations	75
6 Children in Multiple Organizations	76
7 Changing Your User Profile	80
8 Accessing the User Preferences Tab	81
9 Viewing Organizations	87
10 Clicking the Organization Link	91
11 Active Organization	91
12 An Example of Direct Attached Storage	95
13 Selecting the Date and Time from the Calendar	116
14 About the File System Viewer Scan Checker	117
15 Report Views and Report Cache Refresh	119
16 Selecting the Date and Time from the Calendar	124
17 An Example of Global Reporting	131
18 Selecting the Date and Time from the Calendar	140
19 Expanding the Fabric Node	168
20 Fabric Members Are Highlighted	168
21 Highlighting the Applications in the Topology	169
22 Members of a Zone Set	170
23 Displaying Zone Members	171
24 Zone Member to Switch	171
25 Obtaining Information About a Zone Member's Adapter	172
26 Highlighting Elements with Host Bindings	172
27 Displaying Host Bindings	173
28 HBA Port Properties	173
29 WWN Properties	174
30 WWN Properties	174

31	Obtaining Path Information	175
32	Path Information Visible in the Tree	176
33	Enclosing the Elements	186
34	Dragging Multiple Elements to Their New Location	186
35	Showing the Impact of an Application	191
36	Determining Business Cost	193
37	Expanding the Topology Pane	194
38	Filtering Fabrics	194
39	Remote Console	211
40	Obtaining Information About a Host	215
41	Details of a Host Connected to a Switch	216
42	Finding the Status of a Port	218
43	Topology of a Server	224
44	Topology of a Server (Continued)	225
45	Multipathing Displayed in the Topology	226
46	Multipathing Displayed in the Topology (Continued)	227
47	Direct Attached Storage in the Topology	228
48	New Window Option	230
49	Viewing Asset Records	235
50	CXFS File System	238
51	Resources in Two Zones	241
52	Overview of Zoning Capabilities	242
53	Selecting a Storage System	294
54	Selecting a Storage System	301
55	Selecting a Storage System	305
56	Selecting a Storage System	309
57	Specifying a LUN Number	311
58	Deleting Multiple Provisioning Jobs	316
59	Hiding the Middle Pane	342
60	Selecting Organizations Used in This Report	346
61	Selecting Organizations Used in This Report	348
62	Report Architecture	351
63	Choosing a Standard Report	354
64	Adding Tables for a Standard Report	355
65	Linking Common Data in Tables for a Standard Report	356
66	Creating Search Criteria for Standard Reports	357
67	Deciding Which Data Should Appear in the Report	358
68	Sorting Information in the Report	359
69	Selecting the Layout of the Report	360
70	Report Template Displayed	361
71	Clicking the View Tab	362
72	Legend for Capacity Manager	424
73	Capacity of an Element	425
74	Viewing the Capacity of Elements in a Fabric	428
75	Viewing a Utilization Report	431
76	Chargeback Capacity	476
77	Capacity in Capacity Manager	476

Tables

1	Document conventions	xx
2	Status Light Settings	7
3	Discovery Requirements for Switches	14
4	Required Switch Models and InVsn Versions for Discovery16	
5	Discovery Settings for McDATA and Connectrix Switches	19
6	Discovery Requirements for Storage Systems and NAS Filers.	27
7	Default Role Privileges	72
8	Default Role Privileges with Elements	73
9	Changing User Preferences for Event Monitoring for Storage Essentials	81
10	License Requirements	93
11	Determining Managed Access Points	94
12	Logging Levels.	108
13	Default Settings for Clearing Events	112
14	About the Collectors for File System Viewer	114
15	About Collectors for Reports	121
16	Viewing System Wide E-mail Schedules	124
17	Description of "An Example of Global Reporting"	131
18	About Performance Collectors	138
19	Description of Backups.	144
20	Backup Directories in %ORA_HOME%\rman146	
21	Backup Example	147
22	Default Passwords for Database Accounts.	158
23	Description of Toolbar in System Manager	164
24	Elements Displayed in the Topology	167
25	Drop-Down Menu Options Accessible from the Topology*	177
26	Drop-Down Menu Options on the Access Tab	181
27	Show Impact Results.	192
28	Legend for Severity Level Icons	195
29	Descriptions of the Severity Levels Presented in the Severity Menu195	
30	Variables for All Elements*	208
31	Variables for Only Storage Systems, Switches, and Hosts*	208
32	Variables for Only Switches	209
33	Variables for Only Hosts*	209
34	Variables for Only Applications.	210
35	Buttons on the Remote Console	212
36	Drop-Down Menu Options	213
37	Information Available from the Navigation Page	216
38	Port Status Definitions.	218
39	About the Toolbar in the Topology Tab.	229
40	Setting Up Zoning	242
41	Zoning Support	243
42	Provisioning and Pool Support.	257
43	Volume and HSG Support	258

44	Setting Up Storage Partitioning	259
45	HSG Port Assignments	270
46	Allowed Initiators in Host Security Groups	272
47	Volume Usage	288
48	Overview for Path Provisioning	292
49	Zone Icons	297
50	Zone Icons	307
51	Zone Icons	321
52	About the Icons in Event Monitoring for Storage Essentials	324
53	Supported Hardware	324
54	Default Settings for Clearing Events	328
55	Brocade Switch Events	331
56	Severity Levels	333
57	Viewing E-mail Schedules for a Report	349
58	Description of the Report Views	364
59	MVC_HOSTSUMMARYVW	367
60	MVC_CARDSUMMARYVW	368
61	MVC_HOSTVOLUMESUMMARYVW (logical volumes)	369
62	MVC_HOSTDISKDRIVEVW	369
63	MVC_STORAGESYSTEMSUMMARYVW	370
64	MVC_STORGAEPOOLSUMMARYVW	370
65	MVC_STORAGEVOLUMESUMMARYVW	371
66	MVC_SWITCHSUMMARYVW	372
67	MVC_PORTSUMMARYVW	374
68	MVC_ZONESUMMARYVW	374
69	MVC_ZONEVW	375
70	MVC_PATHVW	376
71	MVC_SUBPATHVW	376
72	MVC_MULTIPATHVW	377
73	MVC_EVENTSVW	377
74	MVC_ORGANIZATIONVW	378
75	MVC_ORGRELATIONVW	378
76	MVC_HOSTCAPACITYVW	379
77	MVC_STORAGESYSTEMCONFIGVW	379
78	MVC_STORAGEPOOLCONFIGVW	380
79	MVC_SWITCHCONFIGVW	380
80	MVC_OPTIONALTABLEVW	380
81	MVC_DISKEXTENTSUMMARYVW	381
82	MVC_STORAGEVOLUMEPORTS	381
83	MVC_VOLUMEDISKDRIVEVW	382
84	MVC_STORAGEPROCESSORSUMMARYVW	382
85	MVC_DISKDRIVESUMMARYVW	383
86	MVC_DISK_EXTENTVW	384
87	MVC_ASSETSUMMARY	384
88	MVC_APPLICATIONSUMMARYVW	386
89	MVC_UNITACCESSVW	387
90	MVCA_DBAPPCAPACITYVW	387

91	MVCA_EXCHAPPCAPACITYVW	388
92	MVCA_VIRTUALAPPCAPACITYVW	388
93	MVCA_FSRM_VOLUMESUMMARYVW	389
94	MVCA_FSRM_AGESUMMARYVW	389
95	MVCA_FSRM_EXTDETAILSUMMARYVW	390
96	MVCA_FSRM_DIRDETAILSUMMARYVW	390
97	MVCA_FSRM_USERSUMMARYVW	390
98	MVCA_FSRM_GROUPSUMMARYVW	391
99	MVCA_FSRM_TOPNFILES	391
100	MVCA_FSRM_AGEDFILEDETAILS	392
101	MVCA_FSRM_LARGEDIRINFO	392
102	MVCA_BU_MASTERSERVERSUMMARY	392
103	MVCA_BU_MEDIASERVERSUMMARY	393
104	MVCA_BU_CLIENTSUMMARY	393
105	MVCA_BU_MEDIASUMMARY	394
106	MVCA_BU_JOBSUMMARY	395
107	MVCA_BU_LIBRARYSUMMARY	396
108	MVIEWCORE_STATUS	397
109	MVIEW_STATUS	397
110	Views from Previous Releases	398
111	Toolbar in Lower Pane of Performance Manager	401
112	About the Monitoring Options	408
113	Host Monitoring Support	419
114	Toolbar in Lower Pane of Capacity Manager	425
115	Explanation of the Properties of the Capacity Levels for HDS Array Groups	430
116	Policy Templates	439
117	Policy Table Description	453
118	Setting Up Chargeback Manager	459
119	Element Type Icons	462
120	Viewing E-mail Schedules for a Chargeback Report	490
121	Element Types	492
122	Troubleshooting Firewalls	505
123	Time-out Properties	513
124	Retry Properties	513
125	Known Device Issues	524

About this guide

This guide provides information about:

- Discovering elements
- Managing Security
- Configuring HP
- Monitoring and managing elements
- Provisioning elements
- Viewing events from elements
- Viewing reports about elements
- Viewing performance, capacity and chargeback information
- Troubleshooting HP

Intended audience

This guide is intended for:

- Network Engineers
- Administrators
- Any one that needs to monitor and/or manage their file servers

Prerequisites

Prerequisites for using this product include:

- Networking
- Storage Area Networks (SANs)
- The Common Information Model (CIM)

Related documentation

In addition to this guide, please refer to other documents for this product:





- Online help for HP Storage Essentials 5.0
- HP Storage Essentials 5.0 Integration Guide
- HP Storage Essentials 5.0 Installation Guide
- HP Storage Essentials 5.0 Application Guide
- HP Storage Essentials 5.0 CLI Guide
- HP Storage Essentials 5.0 for File Servers Guide

These and other HP documents can be found on the HP web site: <http://www.hp.com/support/>

Document conventions and symbols

Table 1 Document conventions

Convention	Element
Medium blue text: Figure 1	Cross-reference links and e-mail addresses
Medium blue, underlined text (http://www.hp.com)	Web site addresses
Bold font	<ul style="list-style-type: none">• Key names• Text typed into a GUI element, such as into a box• GUI elements that are clicked or selected, such as menu and list items, buttons, and check boxes
<i>Italics font</i>	Text emphasis
Monospace font	<ul style="list-style-type: none">• File and directory names• System output• Code• Text typed at the command-line
<i>Monospace, italic font</i>	<ul style="list-style-type: none">• Code variables• Command-line variables
Monospace, bold font	Emphasis of file and directory names, system output, code, and text typed at the command line

-  **WARNING!** Indicates that failure to follow directions could result in bodily harm or death.
-  **CAUTION:** Indicates that failure to follow directions could result in damage to equipment or data.
-  **IMPORTANT:** Provides clarifying information or specific instructions.
-  **NOTE:** Provides additional information.

 **TIP:** Provides helpful hints and shortcuts.

HP technical support

Telephone numbers for worldwide technical support are listed on the HP support web site:

<http://www.hp.com/support/>.

Collect the following information before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

For continuous quality improvement, calls may be recorded or monitored.

HP strongly recommends that customers sign up online using the Subscriber's choice web site at

<http://www.hp.com/go/e-updates>.

- Subscribing to this service provides you with e-mail updates on the latest product enhancements, newest versions of drivers, and firmware documentation updates as well as instant access to numerous other product resources.
- After signing up, you can quickly locate your products by selecting **Business support** and then **Storage** under Product Category.

HP-authorized reseller

For the name of your nearest HP-authorized reseller:

- In the United States, call 1-800-345-1518.
- Elsewhere, visit the HP web site: <http://www.hp.com>. Then click **Contact HP** to find locations and telephone numbers.

Helpful web sites

For third-party product information, see the following HP web sites:

- <http://www.hp.com>
- <http://www.hp.com/go/storage>
- <http://www.hp.com/support/>

1 Overview

This chapter describes the following:

- "About This Product" on page 1
- "Suggested Topics for First-Time Users" on page 4
- "Product Components" on page 4
- "About the User Interface" on page 6
- "Installing the Java Plug-in" on page 8
- "Installing the Software Security Certificate" on page 8
- "Stopping the Service for the Management Server" on page 10

About This Product

IMPORTANT: If you access HP Systems Insight Manager through HTTP over SSL (HTTPS), you must provide the full DNS name for the host to be able to access HP Storage Essentials. For example, you could access HP Systems Insight Manager by using `https://mycomputer.domainname.com:50000`, but you could not use `https://mycomputer:50000`. For non-secure connections (HTTP), the full DNS name does not need to be provided.

HP Storage Essentials is a family of open, standards-built SAN management and storage resource management (SRM) solutions that enable you to visualize, monitor, report on, provision, and manage your heterogeneous storage environment with unprecedented simplicity and speed.

HP Storage Essentials is a family of open, standards-built SAN management and storage resource management (SRM) solutions that enable you to visualize, monitor, report on, provision, and manage your heterogeneous storage environment with unprecedented simplicity and speed.

Storage Essentials is integrated throughout HP Systems Insight Manager. You can determine where Storage Essentials begins by looking at the menus. Whenever a menu item displays "Storage Essentials," it is a gateway into Storage Essentials. For example, assume you selected **Tasks & Logs > View Storage Essentials Logs**. You have selected to view the logs for Discovery Data Collection in Storage Essentials, not in HP Systems Insight Manager. When you select **Tools >**

Storage Essentials, you see a submenu listing several features. These features are only in Storage Essentials.

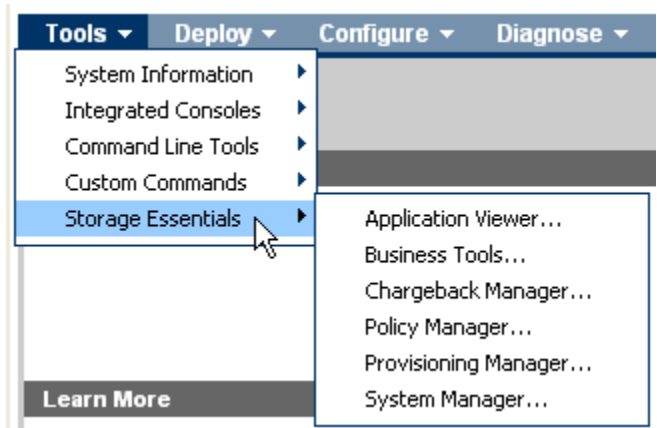


Figure 1 Accessing Features from the Tools > Storage Essentials Menu

Some features are not directly accessible from the menus in HP Systems Insight Manager. For example, to edit a role, you must access Storage Essentials first. You can access Storage Essentials numerous ways. One way to access Storage Essentials is to select **Options > Storage Essentials > Email Settings**. Then, select **Security > Roles** in the upper-right corner of the panel for Storage Essentials.



Figure 2 Accessing Roles in Storage Essentials

To access the online help for Storage Essentials, you must access Storage Essentials as described in the previous paragraph. Then, click **Help > For this page** or **Help > Documentation Center**.

The Documentation Center provides the documentation set for Storage Essentials, including the PDFs, online help, release notes, and support matrix.




Figure 3 Accessing the Documentation Center

To manage and monitor your elements in Storage Essentials, discover your elements in HP Systems Insight Manager. Then, run Discovery Data Collection in Storage Essentials. Discovery Data Collection obtains additional information about your elements. Additional steps are required for managing and monitoring applications and hosts, as mentioned in "[Discovering Applications and Hosts](#)" on page 49.

Storage Essentials runs on a central server, and it relies upon CIM Extensions to provide information about hosts and applications. A CIM Extension is a small software program that gathers information about a host and its applications. The CIM Extension then provides the information it has gathered to Storage Essentials. If you want to monitor a host and its applications, install a CIM Extension on the host.

See "[Discovering Filers, Tape Libraries, Switches, and Storage Systems](#)" on page 13 and "[Discovering Applications and Hosts](#)" on page 49 for more information. Refer to the *Installation Guide* for more information on how to install CIM Extensions.

Storage Essentials can simplify your complex environment and lower your cost of management with CIM-based integrated storage management. The management software integrates the management of applications, servers, storage networks and storage systems in a single, easy to implement and intuitive solution.

The management server integrates the various components in the storage area network infrastructure into a CIM/WBEM/SMI standards based database so you can eliminate any vendor dependencies and view and manage your infrastructure as a whole. A SAN is a network configuration that is dedicated to transporting storage data among network devices, such as storage systems, servers, switches, and tape libraries. Since the SAN is dedicated to transporting storage data, it frees up the data network for regular TCP/IP traffic.

By giving your administrators a single, integrated console to manage tactical activities such as provisioning storage, managing real time events, installing new applications, and migrating servers and storage, as well as strategic activities such as forecasting, planning and cost analysis. The management server's integrated storage management lowers your cost of acquiring and managing a heterogeneous storage environment.

Storage Management Terms

- **CIM** - A common data model of an implementation-neutral schema for describing overall management information in a network/enterprise environment.

- **Web-Based Enterprise Management (WBEM)** - An initiative based on a set of management and Internet standard technologies developed to unify the management of enterprise computing environments.
- **Storage Management Initiative (SMI)** - A SNIA standard for implementing data storage management using CIM.

See the glossary in this user guide for additional definitions.

Key Benefits

- More efficient use of existing assets
- Increased application availability and performance
- Quicker deployment of storage infrastructure and business applications
- Protection of customer flexibility and investments with a standards-based interface

Suggested Topics for First-Time Users

First become familiar with the discovery and Discovery Data Collection in the management server. You must configure the management server to discover the devices on the network so that the management server becomes aware of them. Then, you must Discovery Data Collection so that the management server is aware of the various types of elements on the network.

To learn more about discovery and Discovery Data Collection, see the following topics:

- ["Discovering Filers, Tape Libraries, Switches, and Storage Systems"](#) on page 13
- ["Discovering Applications and Hosts"](#) on page 49

Once you have discovered and obtained details about the devices in the network, begin adding users and adding them to roles and organizations. See ["Managing Security"](#) on page 71 for more information.

Once you are done adding users and roles, use the following list as a guideline for the topics you should learn about:

- ["Provisioning Manager"](#) on page 239
- ["Event Management"](#) on page 323
- ["Running Reports"](#) on page 337
- ["Viewing Performance Data"](#) on page 399

Product Components

This product ships with the following software:

- **Management server** - The management server provides various tools to let you monitor and manage your SAN devices. See ["Management Server Components"](#) on page 5 for more information about these tools.
- **CIM Extensions** - A CIM Extension gathers information from the operating system and host bus adapters. It then makes the information available to the management server. For more information, see the following in the glossary. Refer to the *Installation Guide* for information on how to install the CIM Extensions.

- **(Optional) Module for managing Microsoft Exchange Server** - Management software for Microsoft Exchange lets the administrator actively manage the data storage requirements for Microsoft Exchange.
- **(Optional) Module for managing Oracle Database** - Management software for Oracle reduces database storage cost and improves performance, availability and reliability by assisting the Oracle database administrator in administering Oracle instances, particularly storage.
- **(Optional) Module for managing Microsoft SQL Server** - Management software for Microsoft SQL Server lets the administrator manage and monitor Microsoft SQL Server.
- **(Optional) Module for managing Sybase Adaptive Server Enterprise** - Management software for Sybase lets the administrator actively manage the data requirements for Sybase.

Management Server Components

You may not have access to all of the features described in this section depending on the following:

- The type of license you have. Depending on your license, all features may not be available. See the “List of Features” to determine if you have access to all features. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).
- The role assigned to your user account. For example, users assigned to the Help Desk role by default have access to Application Viewer and Event Monitoring for Storage Essentials, but not to System Manager, Provisioning Manager, Policy Manager, and Reporting.

The following is the full feature set available:

- **Application Viewer** - Application Viewer lets you monitor and display data from applications. You can access Application Viewer by clicking **Tools > Storage Essentials > Application Viewer**.
- **Business Tools** - The management server provides advisors and automators. Advisors provides detailed information for you to make informed decisions about your network, such as non-compliant HBA firmware. Automators let you automate particular tasks. The advisors and automators available to you depend upon your release. See “[Business Tools](#)” on page 495 for more information.
- **Capacity Manager** - Capacity Manager provides a graphical representation of an element's storage capacity in the storage network. See “[Finding an Element's Storage Capacity](#)” on page 423 for more information.
- **Chargeback Manager** - Chargeback Manager helps you manage departmental ownership, track cost and assemble business reports making inquiries, such as audits and inventory reviews, easier. See “[Chargeback Manager](#)” on page 457 for more information.
- **Command Line Interface** - The command line interface (CLI) provides an alternate way for you to manage elements the management server monitors. You can use the CLI commands in scripts to manage your storage.
- **File System Viewer** - File System Viewer does a recursive lookup on the file system and stores the information in an embedded database. File System Viewer can scan files very quickly because of its structure in the database and because it uses a multi-threaded process. More

than one process can be used at a time to scan the files. Refer to the File Servers Guide for more information.

- **Event Monitoring for Storage Essentials** - Storage Essentials provides detailed information about events from an element. To access Event Monitoring for Storage Essentials, click the **Storage Essentials Managed** link in HP Systems Insight Manager. Click the link for the system name. On the HP System Page, click **SE System Properties**. Then, click the **Events** tab. See "[Event Management](#)" on page 323 for more information.
- **Path Provisioning** - Path Provisioning lets you schedule provisioning task, such as creating zones, to run at a later time. See "[Path Provisioning](#)" on page 289 for more information.
- **Performance Manager** - Performance Manager provides a graphical representation of the performance history of an element, such as bytes transmitted per second for a switch. See "[Viewing Performance Data](#)" on page 399 for more information.
- **Policy Manager** - Policy Manager can automatically send an e-mail, generate an event, or run a remote script when an element is being overused or when one of the following occurs:
 - A new element is discovered
 - Successful provisioning occurred
 - An event occurred on one or more specified elementsSee "[Managing Policies](#)" on page 437 for more information.
- **Provisioning Manager** - Provisioning Manager assists you in creating zones, zone sets, and zone aliases, in addition to storage pools, volumes, and host security groups. For more information, see "[Provisioning Manager](#)" on page 239.
- **Reporting** - Reporting provides a variety of detailed reports, such as dependency, event, and utilization reports for discovered elements. See "[Running Reports](#)" on page 337 for more information.
- **System Manager** - System Manager is the gateway to many features that let you view details about the discovered elements. System Manager provides a topology that lets you view how the devices in your network are connected. For more information, see "[Viewing Element Topology and Properties](#)" on page 161.

About the User Interface

The user interface for the management server is split into three panes:

- **Top pane** - It provides access to discovery and configuration features, in addition to features that let you change your login profile. The management server can be configured to display the icons for the management server's utilities.
- **Left pane** - Displays the status light and the buttons for the management server's utilities.
- **Right pane** - Displays the output of a utility. For example, this pane displays the topology.

About the Top Pane

The menus and button in the upper-right corner of the page provide the following functionality.

IMPORTANT: The **Configuration**, **Security**, and **Discovery** buttons only appear if you belong to a role that has System Configuration selected on the Edit Role page (such as the Domain Administrator role). See "[Managing Security](#)" on page 71 for more information.

- **Configuration** - This menu provides the tools for you to manage the management server, such as saving the database. See "[Configuring the Management Server](#)" on page 99.
- **Security** - This menu lets you manage users, organizations, roles and licenses. See "[Managing Security](#)" on page 71.
- **Discovery** - This menu provides the tools for the management server to discover and obtain information from elements in your network. See "[Discovering Filers, Tape Libraries, Switches, and Storage Systems](#)" on page 13 and "[Discovering Applications and Hosts](#)" on page 49.
- **Help** - This menu provides access to the online help and the copyright information.

The links in the upper-left corner let you modify your profile and sign out of the management server.





The status light indicates the status of the management server. It usually displays green under normal conditions, as shown in the following figure:



Figure 4 Status Light

The status light changes color during the following:

Table 2 Status Light Settings

Scenario	Status Light
Normal	
Discovery	
Backup Topology Details	
Include infrastructure details	

NOTE: When the status light is orange or red, you may want to click the text to the left of the light to access discovery logs quickly (**Tasks & Logs > View Storage Essentials Log**).

.For first-time users, see the following topics:

- “[Discovering Filers, Tape Libraries, Switches, and Storage Systems](#)” on page 13
- “[Discovering Applications and Hosts](#)” on page 49
- “[Managing Security](#)” on page 71
- “[About System Manager](#)” on page 161
- “[About Event Monitoring for Storage Essentials](#)” on page 323
- “[About Provisioning Manager](#)” on page 239
- “[About Reporting](#)” on page 337

Installing the Java Plug-in

Java 2 Runtime Environment is required to access several features in the management server, such as System Manager. If your Web browser is running on Sun Solaris, you must manually install the Java plug-in.

To install the Java plug-in:

1. Go to the following URL and download the installation file for the Sun JRE when asked:

`http://<management_server>/appiq/j2re-1_4_2_04-solaris-sparc.sh`
where `<management_server>` is the hostname of the management server.

2. In a terminal window, execute the downloaded file in a directory where you want the JRE installed.

This executable installs the Sun JRE on your computer.

The Java plug-in for your Web browser is available in the following file:

`$JRE_HOME/plugin/sparc/ns610/libjava_oji.so`

where `$JRE_HOME` is the directory containing the JRE installation.

3. In a terminal window, go to the `$HOME/.mozilla/plugins` directory. Create a `plugins` directory if it does not exist.
4. Remove any existing links to the Java plug-in in this directory.
5. Create a symbolic link to the Java plug-in by using the following command:

```
ln -s $JRE_HOME/plugin/sparc/ns610/libjava_oji.so .
```

NOTE: Remember the dot at the end of the command.

6. If you are a root user on the server and you want to make the plug-in available to all users, create a symbolic link in the `plugins` directory under the browser's installation directory, typically `/opt/SUNWns/plugins`.

NOTE: Any existing plug-ins in a user's home directory take precedence over this system-wide plug-in.

7. Restart your Web browser.

Installing the Software Security Certificate

To stop receiving a Security Alert message each time you use the HTTPS logon, install the software security certificate, as described in this section.

- "Installing the Certificate by Using Microsoft Explorer 6.0" on page 9
- "Installing the Certificate by Using Netscape Navigator 7" on page 9
- "Changing the Security Certificate to Match the Name of the Server" on page 10

Keep in mind the following:

- Enter the DNS name of the computer in the URL instead of localhost. If you use `https://localhost` to access the management server, you are shown a "Hostname Mismatch" error.
- For customers to receive a trusted certificate, they need to purchase a certificate from a trusted entity (most browsers have trust relationships set up for Verisign, Entrust, and Thawte, among others.) with the Common Name (CN) set to the name of their management server. Also note that the Common Name in the certificate must match the name in the URL.

Installing the Certificate by Using Microsoft Explorer 6.0

1. Access the management server by typing the following:
`https://machinename`
where `machinename` is the name of the management server.
2. When the security alert message appears, click **OK**.
If you do not want the Web browser to warn you about a secure connection at any Web site, select the **In the future, do not show this warning** option.
3. When you are told there is a problem with the site's security certificate, click the **View Certificate** button.
4. When you are shown the certificate information, click the **Install Certificate** button at the bottom of the screen.
5. When you are shown the Certificate Import Wizard, click **Next** to continue the installation process.
6. Select one of the following:
 - **Automatically select the certificate store based on the type of certificate** - This option places the certificate automatically in the appropriate location.
 - **Place all certificates in the following store** - This option lets you pick the store where the certificate will be stored.
7. Click **Finish**.

8. When you are asked if you want to install the certificate, click **Yes**.
You are shown the following message when the certificate is installed.

Installing the Certificate by Using Netscape Navigator 7

1. Access the management server by typing the following:
`https://machinename`
where `machinename` is the name of the management server.
2. When the security alert message appears, click the **Always** button.
3. When you are told you are requesting an encrypted page, click **OK**.
4. Click the **Always** button when you are asked if you want to accept the certificate.
5. When asked if you wanted to trust the signed applet, click the **Always** button.

Changing the Security Certificate to Match the Name of the Server

If your users are shown a Security Alert window with the following message, you might want to modify the security certificate so users feel more comfortable with installing the certificate:

"The name of the security certificate is invalid or does not match the name of the site."

You can change the security certificate so that users receive the following message instead:

"The security certificate has a valid name matching the name of the page you are trying to view."

When you change the certificate, you must use the `generateAppiqKeystore.bat` program to delete the original certificate. Then, use the `generateAppiqKeystore.bat` program to create a new certificate and to copy the new certificate to the management server.

To change the certificate:

1. Go to the `[Install_Dir]\Tools` directory, where `[Install_Dir]` is the directory into which you installed the management server.
2. To delete the original certificate, enter the following at the command prompt:
`C:\[Install_Dir]\Tools> generateAppiqKeystore.bat del`
The original certificate is deleted.
3. To create a new certificate containing the DNS name of the management server, enter the following at the command prompt:
`C:\[Install_Dir]\Tools> generateAppiqKeystore.bat`
4. If the program is unable to detect a DNS name, enter the following at the command prompt:
`C:\[Install_Dir]\Tools> generateAppiqKeystore.bat mycomputername`
where `mycomputername` is the DNS name of the computer
5. To copy the new certificate to the management server, enter the following at the command prompt:
`C:\[Install_Dir]\Tools> generateAppiqKeystore.bat copy`
The new certificate is copied to the correct location.

Stopping the Service for the Management Server

By default, the software runs as a service from the time the management server has restarted. If you must stop the service, it is recommended that you stop the service from the Services window, rather than from the command prompt window.

IMPORTANT: The service must be running for users to access the management server.

To stop the service on a management server:

1. Go to the **Services** window, usually accessible from the Control Panel.
2. Right-click **AppStorManager**.
3. Select **Stop** from the drop-down menu.
4. To start the management server, right-click **AppStorManager** and select **Start** from the drop-down menu.

2 Discovering Filers, Tape Libraries, Switches, and Storage Systems

Before you can use the management server, you must make the software aware of the elements on your network. An element is anything on the network that can be detected by the management server, such as a switch. This is done through the discovery process. Discovery obtains a list of discovered elements and information about their management interface and dependencies. The management server can discover only elements with a suitable management interface. Refer to the support matrix for supported hardware.

First discover the switches, storage systems, filers, and tape libraries on your network by using the tools in Storage Essentials. Refer to the documentation available from Storage Essentials. Then, run Discovery Data Collection. Discovery Data Collection is required to obtain information from your switches and storage systems. Discovery Data Collection takes some time. You might want to perform this process when the network and the managed elements are not busy. After Discovery Data Collection, install the CIM Extensions. Then, discover your hosts, as described in ["Discovering Applications and Hosts"](#) on page 49.

Discovery of switches, storage systems, filers and tape libraries consists of several steps:

1. Discover your switches by using HP Systems Insight Manager. See ["Step 1 - Discover Switches"](#) on page 14.
2. Discover your storage systems, filers and tape libraries by using HP Systems Insight Manager. See ["Step 2 - Discover Storage Systems, Filers and Tape Libraries"](#) on page 26.
3. Obtain details about the elements you previously discovered by using the Discovery Data Collection feature in Storage Essentials. See ["Step 3 - Discovery Data Collection"](#) on page 39.

To save time, make sure the user names and passwords are correct. The management server tries each of the default user names and passwords whenever it finds an element.

Keep in mind the following:

- After you discover an EMC Connectrix or McDATA switch, the IP address displayed next to the name of the switch is actually the IP address of the service processor for the switch in the Discovery Data Collection screens. To find the IP address of the switch, click the link for the switch in the Discovery Data Collection screen (**Options > Storage Essentials > Discovery > Run Discovery Data Collection**) and then click the Properties tab. The Properties tab can also be accessed by double clicking the switch in System Manager. Complete the steps in this chapter before you try to find the IP address of the switch.
- If you are having a problem with discovering an element, see ["Troubleshooting"](#) on page 501.
- The IP addresses of excluded elements appear in the Discovery Data Collection lists (**Options > Storage Essentials > Discovery > Run Discovery Data Collection**). The management server does not display additional information about excluded elements in the user interface. .
- To obtain information about the storage area network (SAN), include in the discovery the IP addresses for the following:

- **Fibre channel switch** The fibre channel switch contains a list of all elements within the fabric. The management server obtains a detailed listing of all elements connected to the switch fabric.
- **A host containing a Host Bus Adapter (HBA)** All fibre channel host adapters look for available elements attached to the HBA. This information is gathered by CIM Extensions and sent to the management server. Since you have not installed CIM Extensions yet, the management server obtains limited information on the hosts when you perform discovery this time around.
- **A proxy connected to the SAN** - Include a proxy that has a direct connection or a SAN connection to the management server. An example of a proxy is the EMC Solutions Enabler or Hitachi HiCommand Device Manager. Engenio storage systems do not require a proxy, as they can be accessed directly. Make sure the proxy service has started. On a computer running Windows, this can be determined by looking in the Services window. EMC Solutions Enabler version 5.1 requires additional steps for discovery. See "[Discovering EMC Solutions Enabler 5.1](#)" on page 28 for more information.

Step 1 - Discover Switches

This section describes the following:

- "[Discovering Brocade Switches](#)" on page 15
- "[Discovering CNT Switches](#)" on page 16
- "[Discovering Cisco Switches](#)" on page 17
- "[Discovering Sun StorEdge and QLogic Switches](#)" on page 17
- "[Changing the SNMP Trap Listener Port for Sun StorEdge Switches](#)" on page 18
- "[Discovering McDATA and EMC Connectrix Switches](#)" on page 18
- "[Excluding McDATA and EMC Connectrix Switches from Discovery](#)" on page 24

The following table provides an overview of the discovery requirements for switches.

Table 3 Discovery Requirements for Switches

Element	Discovery Requirements	Additional Information
Brocade switches	Enter the IP address/DNS name, user name and password of the Brocade switch to discover it. The user name (default admin) and password must be the Admin Account.	See " Discovering Brocade Switches " on page 15.
CNT Switches	Enter the IP address followed by the port number for the InVsn Software that manages the switch.	See " Discovering CNT Switches " on page 16.

Table 3 Discovery Requirements for Switches (continued)

Element	Discovery Requirements	Additional Information
Cisco Switches	Enter the IP address/DNS name of the Cisco switch. You do not need to enter a password.	See "Discovering Cisco Switches" on page 17.
Sun StorEdge and QLogic switches	Enter the IP address/DNS name of the Sun StorEdge or QLogic switch. You do not need to enter a password.	See "Discovering Sun StorEdge and QLogic Switches" on page 17.
McDATA and EMC Connectrix switches	Additional steps are required for discovering these switches, and the steps vary according to your network configuration.	See "Discovering McDATA and EMC Connectrix Switches" on page 18.

IMPORTANT: Make sure you do not have pop-up blocking software enabled. If your Web browser has an option for blocking pop-ups, disable it. The management server uses pop-ups for dialog boxes.

Discovering Brocade Switches

IMPORTANT: Verify that the Rapid program on the switch is set to 1. Rapid must be set to 1 so that the management server can communicate with the switch. See ["Verifying Brocade Rapid Program Is Set to 1"](#) on page 15 for more information.

To discover Brocade switches, provide the following information in HP Insight Manager:

- IP address or DNS name of the Brocade switch you want to discover.
- User Name for the switch
- Password for the switch

Verifying Brocade Rapid Program Is Set to 1

If you are discovering Brocade switches, verify that the Rapid program on the switch is set to 1.

1. (Optional) Set the command prompt window so that it displays many rows.

While completing the following steps, the command prompt window displays a large amount of data. You might want to expand the size and buffer of the command prompt window. To do this in Microsoft Windows 2000, click the upper-right corner of the command prompt window, click the **Layout** tab, and then modify the options under Screen Buffer Size and Window Size.

2. Access the Brocade switch by using the telnet option. For example,

```
telnet
open 10.1.213.228
```

where 10.1.213.228 is the IP address of the switch.

3. When prompted for the user name and password, supply them.

4. Type the following to see what is supported on the switch:

```
supportshow
```

A large amount of output is displayed.

5. Select all of the output.

6. Paste the output in a text editor, for example Notepad. Use the Find command to search for `rpc.rapid`.

7. Verify Rapid is set to one, as displayed below:

```
rpc.rapid:      1
```

Discovering CNT Switches

The management server uses the CNT SMI-S provider to discover CNT switches. A provider is a small software program that is used by the management server to communicate with a device, such as a switch.

This provider communicates with CNT InVsn Enterprise Manager to obtain information about the switch. The provider requires a certain version of InVsn depending on the switch model. See the following table for more information.

Table 4 Required Switch Models and InVsn Versions for Discovery

Switch Model	InVsn Software Version
FC/9000	9.0 or later
UMD	9.5 or later

Keep in mind the following for CNT switches:

- SNMP is not supported for CNT switches.
- CNT InVsn Enterprise Manager must be running for the management server to discover it.
- The management server does not support provisioning for CNT switches. Only the active zone set and its zone members are reported.
- No ports are reported for uninstalled blades or GBICs.

To discover CNT switches:

1. Before you can discover a CNT switch, you must do the following in the CNT InVsn Enterprise Manager software:

a. Open the file `ProductInfo.ini` in a text editor, such as Notepad. If the software was installed in the default directory, this file should be in the following directory:

```
\Program Files\CNT\inVSN_EM
```

b. Change the following entry in the file:

```
cimomenabled=TRUE
```

- c. Save the file, then restart the InVsn software.
2. In the **IP Address/DNS Name** field in HP Insight Manager, type the primary IP address of the host running the InVsn software you want to discover followed by its port number. For example, if the InVsn software is running at 192.168.10.76 on port 5989, you would specify the IP Address and port number as follows:
192.168.10.76:5989
3. In the **User Name** field, type the user name for the login to the InVsn software.
4. In the **Password** field type the password for the login to the InVsn software.
5. In the **Verify Password** field type the password you provided previously.
6. Start discovery.

Discovering Cisco Switches

The management server discovers Cisco switches through an SNMP connection. When you discover a Cisco switch, you do not need to provide a password.

Keep in mind the following for Cisco switches:

- You can view zones, zone sets and zone aliases on a Cisco switch; however, you cannot use the management server to create, modify or remove them from a Cisco switch.
- The management server gathers information about the Cisco inactive database during Discovery Data Collection. You can change the amount of information that is collected by modifying a property. See ["Changing the Amount of Information Collected from the Inactive Zone Database \(Cisco Switches\)"](#) on page 255.
- The management server groups active zone sets in all Virtual SANs (VSANs) in a fabric into a zone set called "ACTIVE", and the "ACTIVE" zone set is shown associated with the physical fabric. The members of the "ACTIVE" zone set (zones, zone sets, zone aliases) have the name of the VSAN prefixed to their name. For example, an active zone named "ZONE1" from a VSAN named "VSAN1" is displayed as a zone on the physical fabric with name "VSAN1:ZONE1".

To discover a Cisco switch in HP Insight Manager, you must provide the following:

- DNS name or primary IP address of the Cisco switch you want to discover
- User name for the switch. The password can be left blank.

Discovering Sun StorEdge and QLogic Switches

The management server discovers Sun StorEdge and QLogic switches through an SNMP connection. When you discover a Sun StorEdge or QLogic switch, you do not need to provide a password.

Keep in mind the following:

- The management server does not support provisioning for Sun StorEdge and QLogic switches. Only the active zone set and its zone members are reported.
- To manage a fabric of Sun StorEdge and/or QLogic switches, every switch in the fabric must be included in the discovery list. If a switch is not included in the discovery list, it may show up as a generic host system.

- No ports are reported for uninstalled blades or GBICs.
- The default SNMP trap listener port for all Sun StorEdge switches is 162. To change this port, see [“Changing the SNMP Trap Listener Port for Sun StorEdge Switches”](#) on page 18.
- To receive events from Sun StorEdge switches, verify the SNMP trap community string is set to public in SANbox Manager or via telnet. Also, make sure the SNMP traps are configured to be sent to the management server. See [“About Event Monitoring for Storage Essentials”](#) on page 323 for more information.

To discover a Sun StorEdge or QLogic switch in HP Insight Manager, provide the following:

- IP Address or DNS name of the Sun StorEdge or QLogic switch you want to discover.
- The user name for the switch. This is the public community SNMP string (read community password) for the switch. This field can be left blank if the element's user name and password are one of the default user names and passwords. You do not need to provide a password.

Changing the SNMP Trap Listener Port for Sun StorEdge Switches

The default SNMP trap listener port for all Sun StorEdge switches is 162. To change this port for all switches that are discovered through SNMP, modify the `cimom.snmpTrapListenerPort` property as described in the following steps:

1. Select **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the `cimom.snmpTrapListenerPort` property. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
4. Return to the Advanced page (**Options > Storage Essentials > Manage Product Health**). Then, click **Advanced** in the Disk Space tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make your changes in the **Custom Properties** field. Make sure the property is not commented out by removing the hash (#) symbol in front of the property.
7. Set the `cimom.snmpTrapListenerPort` property to the port you want, as shown in the following example:

```
cimom.snmpTrapListenerPort=162
```

8. When you are done, click **Save**.
9. Restart the service for the management server for your changes to take effect.
While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Discovering McDATA and EMC Connectrix Switches

McDATA and EMC Connectrix switches use the Fibre Channel Switch Application Programming Interface (SWAPI) to communicate with devices on the network. The management server can discover multiple instances of Enterprise Fabric Connectivity Manager.

Use one of the following techniques to discover McDATA and Connectrix switches:

Table 5 Discovery Settings for McDATA and Connectrix Switches

Discovery	SWAPI setting through a Proxy	SNMP setting Through a Proxy	Contacting the switch directly
Description	Use this option if you have Enterprise Fabric Connectivity (EFC) Manager. You will need to connect through the proxy instead of the switch. See "SWAPI Setting Through a Proxy" on page 20 for more information.	Contact the switch through a proxy. You can use this option with EMC Connectrix™ Manager and Enterprise Fabric Connectivity (EFC) Manager to contact the switch. See "SNMP Setting Through a Proxy" on page 22 for more information.	Contact the switch by its IP address or DNS name. This connection uses SNMP. See "Contacting a McDATA or Connectrix Switch Directly" on page 23.
Provisioning Limitations	The SWAPI setting lets you activate a zone set, in addition to creating, editing, and deleting zones and zone sets. You cannot manage or view information about zone aliases.	This SNMP setting through a proxy does not let you manage or access information about zones, zone sets or zone aliases.	This SNMP setting provides view only access to the active zone set and its members. You cannot create, modify, and/or delete zone sets or its members.

Keep in mind the following:

- If you change a discovery configuration from SNMP to SWAPI or vice versa, the user ID and password will no longer work. For this reason, it is recommended that you set this property before discovering any McDATA switches. If you must change the configuration, see ["Changing the Discovery Settings"](#) on page 24.
- After you discover a McDATA or Connectrix switch through a proxy, the IP address displayed next to the name of the switch is actually the IP address of the proxy for the switch in the Discovery, and Discovery Data Collection screens. To find the IP address of the switch, click the link for the switch in the Discovery Data Collection screen (**Discovery > Details**) and then click the **Properties** tab. The **Properties** tab can also be accessed by double clicking the switch in System Manager.
- If you want to add, remove, or replace McDATA or Connectrix switches after you have discovered the service processor, you must perform additional steps, see ["About Managing McDATA and EMC Connectrix Switches"](#) on page 43.

- If you have problems obtaining information from McDATA or Connectrix switches during discovery and/or Discovery Data Collection, see ["Step 2 - Discover Storage Systems, Filers and Tape Libraries"](#) on page 26.
- All McDATA switches in a fabric must be managed by the same EFC Manager. Do not have more than one EFC Manager to a fabric for McDATA switches. If you do use more than one EFC Manager in a fabric, you must use the same EFC Manager for your zoning. Do not use the other EFC Managers for zoning, as this will create zoning database problems.
- All Connectrix switches in a fabric must be managed by the same Connectrix Manager. Do not have more than one Connectrix Manager to a fabric for Connectrix switches. If you do use more than one Connectrix Manager in a fabric, you must use the same Connectrix Manager for your zoning. Do not use the other Connectrix Managers for zoning, as this will create zoning database problems.
- If you want the management server to receive SNMP events from Connectrix or McDATA switches, do one of the following:
 - If you discovered Connectrix Manager or EFC Manager, only enable SNMP trap forwarding to the management server on the Connectrix Manager or EFC Manager, not on the individual switches. Connectrix Manager or EFC Manager should be configured to forward SNMP traps to the IP address of the management server, and the community string should match the user ID you used to discover Connectrix Manager or EFC Manager.
 - If you discovered Connectrix or McDATA switches directly, enable SNMP trap forwarding on the switches, not on any other management software. The switches should be configured to forward SNMP traps to the IP address of the management server, and the community string should match the user ID you used to discover the Connectrix or McDATA switches.

SWAPI Setting Through a Proxy

With the SWAPI setting, the management server contacts a proxy to obtain information about the switches connected to it. Use Enterprise Fabric Connectivity (EFC) Manager for this option. If you do not have EFC Manager, see ["SNMP Setting Through a Proxy"](#) on page 22. EFC Manager versions 7.0 and later can communicate with the management server and the switch. EFC Manager accesses the switch through a SWAPI connection. This configuration lets multiple instances of the management server or other clients contact EFC Manager, which in turn provides information about the switch.

IMPORTANT: This option only supports McDATA switches. If you want to discover EMC Connectrix switches, you must discover them through the SNMP provider, either directly or through a proxy. See ["SNMP Setting Through a Proxy"](#) on page 22 for more information about using the SNMP provider to discover switches through a proxy. See ["Contacting a McDATA or Connectrix Switch Directly"](#) on page 23 for more information discovering switches by their IP address.

Step 1 - (McDATA Switches Only) Install the Bridge Agent

To communicate with EFC Manager, the management server requires the Bridge Agent. Refer to your McDATA representative for more information about the Bridge Agent.

Step 2 - Change the Discovery Setting for McDATA and Connectrix Switches to SWAPI

To change the discovery settings to SWAPI:

1. Click **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following property. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.

```
cimom.useSnmppMcDataProvider=FALSE
```

4. Return to the Advanced page (**Options > Storage Essentials > Manage Product Health**). Then, click **Advanced** in the Disk Space tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Uncomment the `cimom.useSnmppMcDataProvider` property by removing the number sign (#) in front of `cimom.useSnmppMcDataProvider`.
7. Verify the `cimom.useSnmppMcDataProvider` property is set to false.

```
cimom.useSnmppMcDataProvider=FALSE
```
8. When you are done, click **Save**.
9. Restart the service for the management server for your changes to take effect.
While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Step 3 - Discover the Proxy

To discover the proxy, you must provide the following information to HP Insight Manager:

- IP address or DNS name of the EFC Manager/Connectrix Manager you want to discover
- User name - Type the user name for EFC Manager/Connectrix Manager.
This field can be left blank if one or more of the following conditions are fulfilled:
 - The element's user name and password are one of the default user names and passwords.
 - The element does not require authentication.
- Password - Type the corresponding password for EFC Manager/Connectrix Manager.
This field can be left blank if one or more of the following conditions are fulfilled:
 - The element's user name and password are one of the default user names and passwords.
 - The element does not require authentication.

CIM_ERR_FAILED When Trying to Activate a Zone Set Using McDATA SWAPI

When the user tries to activate a zone set using McDATA SWAPI, the operation may return CIM_ERR_FAILED with one of the following detailed messages:

```
Cannot activate zone set. SWAPI Handle is not valid for fabric  
Cannot activate zone set. Active zone set information is out of date for  
fabric
```

```
There is no active SWAPI connection for fabric
Fabric is not in the cache
```

These error messages indicate that the SWAPI connection to the EFCM managing the fabric is no longer valid, or the active zone information was changed on the fabric without using the management server. The management server does not activate a zone set under these conditions.

To fix this problem, re-discover the EFCM to re-establish the SWAPI connection.

Once the connection is working, the provisioning operation should succeed. If it continues to fail because the active zone set information is out of date, do a Discovery Data Collection for this element to update the zoning information.

SNMP Setting Through a Proxy

This SNMP setting through a proxy does not let you manage or access information about zones, zone sets or zone aliases.

This option is required if you want to discover McDATA or Connectrix switches through a proxy using the SNMP provider. You can use this option with EMC Connectrix™ Manager and Enterprise Fabric Connectivity (EFC) Manager to contact the switch.

Step 1 - Verify the Discovery Setting for Switches Is Set to SNMP

By default the discovery settings for McDATA and Connectrix switches is set to SNMP. If you believe it has been changed to SWAPI, you can perform the following steps to change it back to SNMP.

1. Click **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
2. Click the **Edit** button at the bottom of the page.
3. Comment out the `cimom.useSnmppMcDataProvider` property by placing the number sign (#) in front of `cimom.useSnmppMcDataProvider`.
`#cimom.useSnmppMcDataProvider=false`
4. Restart the service for the management server.
5. Verify the following on the proxy and the switches accessible from the proxy:
 - The SNMP agent is enabled.
 - The read-only community string is configured.

Step 2 - Discover the Proxy

To discover the proxy, you must provide the following information to HP Insight Manager:

- IP address or DNS name of the EFC Manager/Connectrix Manager you want to discover
- User name - The default user name, which is "public" (the read-only community string). This is the user name of the proxy.

IMPORTANT: To access a Windows-based device, prepend the user name with the Windows domain name, as shown in the following example.

```
domain_name\user_name
```

where

- `domain_name` is the domain name of the machine
- `user_name` is the name of your network account
- Password - You can leave the password field blank, where it is being accessed through SNMP.

Step 3 - Make Sure There Are No Port Conflicts for Receiving SNMP Traps

When the management server is configured to contact the proxy by SNMP, the management server receives events from the proxy in the form of SNMP traps. By default, the management server uses port 162 to receive SNMP traps. If another software package is using that port, the management server is unable to receive the traps. To change the port the management server uses:

1. Select **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
2. Click the **Edit** button at the bottom of the screen.
3. Set the `cimom.snmpTrapListenerPort` to another port, as shown in the following example:

```
cimom.snmpTrapListenerPort=1234
```

where 1234 is the new port

4. When you are done making your changes, click the **OK** button.
5. Restart the service for the management server for your changes to take effect.

While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Step 4 - Step Up the Proxy to Send Traps to the Correct Port

When you are using the SNMP setting to discover a proxy, you must configure the SNMP agent on the proxy manager to send traps to the management server using the port you selected. This configuration sends traps from all switches managed by that proxy. Refer to your documentation for your proxy for more information.

Contacting a McDATA or Connectrix Switch Directly

To discover a McDATA or Connectrix switch directly, provide the following to HP Insight Manager:

- The IP address or DNS name of the switch you want to discover.
- The user name for accessing the switch. The default user name is "public" (the read-only community string).
- No password. The password does not matter since the management server is not doing any configurations through SNMP.

Make Sure There Are No Port Conflicts for Receiving SNMP Traps

When the management server is configured to contact a switch by SNMP, the management server receives events from the switch in the form of SNMP traps. By default, the management server uses port 162 to receive SNMP traps. If another software package is using that port, the management server is unable to receive the traps. To change the port the management server uses:

1. Select **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
2. Click the **Edit** button at the bottom of the screen.
3. Set the `cimom.snmpTrapListenerPort` to another port, as shown in the following example:

```
cimom.snmpTrapListenerPort=1234
```

where 1234 is the new port

4. When you are done making your changes, click the **OK** button.
5. Restart the service for the management server for your changes to take effect.
While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Configure the SNMP Agent to Send Traps to the Correct Port

When you are using the SNMP setting to discover a switch, you must configure the SNMP agent on the switch to send traps to the management server using the port you selected. Refer to your documentation for your switch for more information.

Changing the Discovery Settings

To change the discovery settings from SWAPI to SNMP or vice versa:

1. Delete all McDATA and Connectrix switches in the application.
2. Click **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
3. Click the **Edit** button at the bottom of the page.
4. Change the `cimom.useSnmpMcDataProvider` property as follows:
 - **SNMP setting** - Comment out the `cimom.useSnmpMcDataProvider` property by placing a number sign (#) in front of the `cimom.useSnmpMcDataProvider` property as follows: `#cimom.useSnmpMcDataProvider=false`
 - **SWAPI setting** - Remove the number sign (#) in front of the `cimom.useSnmpMcDataProvider` property.
5. Restart the service for the management server.
6. Add new elements in the Discovery screen.
 - **SWAPI connection** - Enter the IP address, user name and password for the proxy.
 - **SNMP connection** - Enter the IP address of the proxy. The default user name is "public" (the read-only community string). The password does not matter since the management server is not doing any configurations through SNMP.
7. Verify the following on the proxy and the switches accessible from the proxy:
 - The SNMP agent is enabled.
 - The read-only community string is configured.
8. Start discovery in HP Insight Manager.
9. Perform Discovery Data Collection.

Excluding McDATA and EMC Connectrix Switches from Discovery

Specific McDATA and Connectrix switches can be excluded from discovery by using system properties.

To exclude one or more switches from discovery, you must modify the `cimom.mcddata.exclude` property. Set the property `cimom.mcddata.exclude` to a comma separated list of Worldwide Names of the McDATA and Connectrix switches you want excluded, as shown in the following example:

```
cimom.mcddata.exclude=1000080088A07024,1000080088A0D0B6
```

The management server excludes the switches with one of the following Worldwide Names: 1000080088A07024 and 1000080088A0D0B6

If the `cimom.mcddata.exclude` property is not modified, the management server discovers and obtains details from all McDATA and Connectrix switches.

IMPORTANT: The IP addresses of excluded elements appear in the Discovery Data Collection lists. The management server does not display additional information about excluded elements in the user interface.

To modify the `cimom.mcddata.exclude` property:

1. Select **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the `cimom.mcddata.exclude` property. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
4. Return to the Advanced page (**Options > Storage Essentials > Manage Product Health**). Then, click **Advanced** in the Disk Space tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make your changes in the **Custom Properties** field. Make sure the property is not commented out by removing the hash (#) symbol in front of the property.
7. Add the Worldwide Names corresponding to the switches you want to exclude from discovery. Separate additional Worldwide Names with a comma, as shown by the following example:

```
cimom.mcddata.exclude=1000080088A07024,1000080088A0D0B6
```

where 1000080088A07024 and 1000080088A0D0B6 are the Worldwide Names for McDATA and Connectrix switches.

8. When you are done, click **Save**.

While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Viewing Log Messages

Use the **View Logs** tab to obtain the status of the following:

- Discovery Data Collection

During these operations, the management server displays its status at regular intervals.

To view logs for these operations:

1. Click the **Tasks & Logs > View Storage Essentials Log**.
2. To obtain the latest status, click the **Get Latest Messages** button.

If the software is unable to discover or obtain information about a device, the log messages might provide some information as to where the problem occurred.

For example, if a host was not discovered, the log messages might indicate the provider configuration for that device was never created. This could mean the software was given the wrong user name and/or password for that host. As a result, the software logged onto the host with a guest account, which does not have enough permissions to start WML.

IMPORTANT: Look at Event Monitoring for Storage Essentials for additional information. See ["About Event Monitoring for Storage Essentials"](#) on page 323 for more information.

Duplicate Logs for Brocade Switches in Same Fabric

If you discover more than one Brocade switch in the same fabric, the discovery log displays duplicate listings for the Brocade switches. Each Brocade switch is listed multiple times with the IP address of the other switches and its own.

For example, assume you are discovering Brocade switches QBrocade2 and QBrocade5 in the same fabric, two duplicate entries are displayed in the log. QBrocade2 is listed twice, once with its own IP address, the other time with the IP address of QBrocade5, as shown below.

```
[Nov 27, 2002 8:45:05 AM] Discovered Switch: QBrocade2 at 192.168.10.22
[Nov 27, 2002 8:45:09 AM] Discovered Switch: QBrocade5 at 192.168.10.22
[Nov 27, 2002 8:45:09 AM] Enabling provider configuration:
APPIQ_BrocadeElementManagerConfig
[...]
[Nov 27, 2002 8:45:37 AM] Discovered Switch: QBrocade2 at 192.168.10.25
[Nov 27, 2002 8:45:42 AM] Discovered Switch: QBrocade5 at 192.168.10.25
[Nov 27, 2002 8:45:42 AM] Enabling provider configuration:
APPIQ_BrocadeElementManagerConfig
192.168.10.22 Switch QBrocade2, QBrocade5 admin
192.168.10.25 Switch QBrocade2, QBrocade5 admin
```

Step 2 - Discover Storage Systems, Filers and Tape Libraries

- ["Discovering EMC Solutions Enabler 5.1"](#) on page 28
- ["Excluding EMC Symmetrix Storage Systems from Discovery"](#) on page 29
- ["Discovering EMC CLARiiON Storage Systems"](#) on page 30

- ["Discovering HDS Storage Systems"](#) on page 31
- ["Excluding HDS Storage Systems from Discovery"](#) on page 32
- ["Discovering HP StorageWorks Arrays"](#) on page 33
- ["Discovering Engenio Storage Systems"](#) on page 33
- ["Discovering NetApp Filers"](#) on page 34
- ["Discovering Sun StorEdge 3510 Storage Systems"](#) on page 35
- ["Discovering Sun StorEdge 6920 Storage Systems"](#) on page 36
- ["Discovering Sun StorEdge 6130 Storage Systems"](#) on page 36
- ["Discovering IBM Storage Systems"](#) on page 36
- ["Discovering IBM Tape Libraries"](#) on page 37

Table 6 Discovery Requirements for Storage Systems and NAS Filers

Element	Discovery Requirements	Additional Information
EMC CLARiiON storage systems	The EMC Navisphere CLI is required for the management server to communicate with the CLARiiON storage system.	See "Discovering EMC CLARiiON Storage Systems" on page 30 for more information.
EMC Symmetrix storage system (Including EMC Symmetrix DMX storage systems)	Discover the server running the EMC Solutions Enabler.	See "Discovering EMC Solutions Enabler 5.1" on page 28 for more information.
HDS storage systems	Discover the server running HiCommand Device Manager.	See "Discovering HDS Storage Systems" on page 31 for more information.
HP storage systems	Discover the server running the HP CIMOM.	See "Discovering HP StorageWorks Arrays" on page 33.

Table 6 Discovery Requirements for Storage Systems and NAS Filers (continued)

Element	Discovery Requirements	Additional Information
Engenio storage systems	<p>Can be discovered two ways:</p> <ul style="list-style-type: none"> • Entering the IP address/DNS name, user name and password of a controller for an Engenio storage system. Discovers only the corresponding IP address of the controller. • Entering the IP address/DNS name, user name and password of a proxy that is used to manage an Engenio storage system. Discovers all controllers known to the proxy. 	See "Discovering Engenio Storage Systems" on page 33.
NetApp Filer	Discover the filer directly.	See "Discovering NetApp Filers" on page 34.
Sun StorEdge 3510	Discovered through proxy software called Sun StorEdge™ Configuration Service. On the discovery page the user should enter the hostname or IP address of the computer running the Sun StorEdge 3510 SMI-S provider.	See "Discovering Sun StorEdge 3510 Storage Systems" on page 35.
Sun StorEdge 6920	Discover the storage system directly.	See "Discovering Sun StorEdge 6920 Storage Systems" on page 36.
Sun StorEdge 6130	Discover the storage system directly. The username does not matter. The password matters only for provisioning.	See "Discovering Sun StorEdge 6130 Storage Systems" on page 36.
IBM Storage Systems	Discover the CIMOM that talks to the IBM storage systems you want to monitor.	See "Discovering IBM Storage Systems" on page 36.
IBM Tape Libraries	Provide the IP address, namespace, user name and password for the tape library.	See "Discovering IBM Tape Libraries" on page 37

Discovering EMC Solutions Enabler 5.1

EMC Solutions Enabler restricts access to itself through the nethost file. If present, the nethost file is located in the same directory as the netcnfg file. If you are using a nethost file, edit it to allow the management server to discover the Solutions Enabler and the Symmetrix storage systems that it manages.

IMPORTANT: Use a nethost file unless you are running a version of the Solutions Enabler earlier than the 5.1 version. You must have the license installed for the Solutions Enabler. The nethost file provides access to the Solutions Enabler API.

Sometimes you can access an EMC Symmetrix storage system through several Solutions Enabler servers. In this case if you do not have access to a particular Solutions Enabler, you may still be able to access the Symmetrix storage system through another Solutions Enabler.

If you do not have a nethost file, you may need to create one. For example, assume you are running Solutions Enabler on a Solaris server, you would create a nethost file as described in the following steps. Refer to the documentation for Solutions Enabler for other operating systems.

1. Create a file called "nethost" in the `/opt/emc/API/symapi/config` directory.
2. Add the following lines to the nethost file:

```
<management server name> SYSTEM  
<management server IP> SYSTEM
```

where
 - `<management server name>` is the DNS name of the management server
 - `<management server name>` is the IP address of the management server
3. Add the following line to the `/opt/emc/API/symapi/config/netcnfg` file:

```
SYMAPI_SERVER - TCPIP <IP of SymAPI server> 2707
```
4. Use the following command to start the daemon:

```
/opt/emc/SYMCLI/V5.5.0/bin/symapisrv -service SYMAPI_SERVER start  
-background
```
5. Use the following command to stop the daemon:

```
/opt/emc/SYMCLI/V5.5.0/bin/symapisrv stop
```
6. You may need to discover the Symmetrix arrays the SymAPI server can see by running the following command:

```
/opt/emc/SYMCLI/V5.5.0/bin/symcfg discover
```

IMPORTANT: If error 214 is present in the discovery log and/or cimom.log during discovery, this means the SymAPI server is not licensed for remote connections. The end-user will have to acquire and install the license before discovery can occur.

Required Licenses

If you want to use all of the features of the management server, such as provisioning, with an EMC Symmetrix storage system, you must have licenses for the following products:

- BASE
- DeltaMark
- SERVER
- DevMasking
- Config Manager
- Mapping (SOLUTION_4)

Excluding EMC Symmetrix Storage Systems from Discovery

When multiple EMC Symmetrix storage systems are managed through a single Solutions Enabler, specific storage systems may be excluded from discovery by using system properties.

To exclude one or more Symmetrix storage systems from discovery, you must modify the `cimom.symmetrix.exclude` property. Set the property `cimom.symmetrix.exclude` to a comma separated list of serial numbers of the storage systems you want excluded, as shown in the following example:

```
cimom.symmetrix.exclude=000183500570,000183610580
```

The management server excludes the storage systems with one of the following serial numbers: 000183500570 and 000183610580.

If the `cimom.symmetrix.exclude` property, the management server discovers and obtains details from all EMC Symmetrix Storage Systems managed by discovered Solutions Enablers.

IMPORTANT: The IP addresses of excluded elements appear in the Discovery Data Collection lists (**Discovery > Details**). The management server does not display additional information about excluded elements in the user interface. The management server, however, does mention in the logs (**Discovery > View Logs**) that a provider instance has been created for an excluded element. You can ignore this message that appears in the logs.

To modify the `cimom.symmetrix.exclude` property:

1. Select **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.

```
#cimom.symmetrix.exclude=000183500570,000183500575
```
4. Return to the Advanced page.
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make sure the property is not commented out by removing the hash (#) symbol in front of the property. Add the serial numbers corresponding to the Symmetrix storage systems you want to

exclude from discovery. Separate additional serial numbers with a comma, as shown by the following example:

```
cimom.symmetrix.exclude=000183500570,000183500575
```

where 000183500570 and 000183500575 are serial numbers for Symmetrix storage systems.

7. When you are done, click **Save**.
8. Restart the service for the management server for your changes to take effect.
While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Discovering EMC CLARiiON Storage Systems

The EMC Navisphere® CLI must be installed on the management server for the management server to communicate with the CLARiiON® storage system. At the time this documentation was created, EMC distributed the Navisphere CLI as part of the EMC Navisphere Software Suite. Contact your EMC representative for more information about obtaining the Navisphere CLI. Distribution rights for the Navisphere CLI belong to EMC. For Solaris, you must install the Navisphere Disk Array Management Tool CLI (NAVICLI).

Contact your EMC representative for more information about obtaining the Navisphere CLI. Distribution rights for the Navisphere CLI belong to EMC.

IMPORTANT: Before you discover your CLARiiON storage systems, you must have already installed all required software components for your CLARiiON storage system, such as the Navisphere Host Agent. Refer to the documentation for your storage system for more information.

In Navisphere Manager add one of the following to the privilege user section:

```
SYSTEM@name_of_my_management_server  
SYSTEM@IP_of_my_management_server
```

where

- `name_of_my_management_server` is the DNS name of the computer running the management server software
- `IP_of_my_management_server` is the IP address of the computer running the management server software

When you use the management server to discover the CLARiiON storage system, provide the IP address for the CLARiiON storage system and the user name and password used to log into Navisphere.

Discovering HDS Storage Systems

HiCommand Device Manager is required for the management server to communicate with an HDS storage system. To discover an HDS storage system, enter the IP address, user name and password

for the server running HiCommand Device Manager. Do not point to the disk array for the storage system.

To obtain information about HDS storage systems, the management server must be able to access the port HiCommand Device Manager uses to listen. By default, HiCommand Device Manager listens on port 2001, and the management server assumes this configuration at discovery time. If HiCommand Device Manager uses a different port, specify this other port when you discover HiCommand Device Manager.

Keep in mind the following:

- You cannot scan an IP range to discover an instance of HiCommand Device Manager that listens on a port other than 2001. The management server does not allow port numbers in the scanning of IP ranges, and thus, you are not able to specify the port.
- The management server communicates with HiCommand Device Manager through a nonsecure connection. If you want the management server to communicate with HiCommand Device Manager through a secure sockets layer (SSL) connection, you must modify an internal property or use HTTPS when you discover HiCommand Device Manager. See ["Communicating with HiCommand Device Manager Over SSL"](#) on page 519.

To discover an HDS storage system that listens on a port other than 2001, you must provide the following information to HP Insight Manager:

- The name of the server and the port HiCommand Device Manager uses to listen separated by a colon, as shown in the following example:

```
proxy2:1234
```

where

- `proxy2` is the name of the server running HiCommand Device Manager
- `1234` is the port HiCommand Device Manager uses to listen
- User name for HiCommand Device Manager.
- Password for HiCommand Device Manager.

Excluding HDS Storage Systems from Discovery

When multiple HDS storage systems are managed through a single HiCommand Device Manager, specific storage systems may be excluded from discovery by using system properties.

To exclude one or more HDS storage systems from discovery, you must modify the `cimom.hds.exclude` property. Set the property `cimom.hds.exclude` to a comma separated list of serial numbers of the storage systems you want excluded, as shown in the following example:

```
cimom.hds.exclude=61038,61037
```

The management server excludes the storage systems with one of the following serial numbers: 61038 and 61037.

If the `cimom.hds.exclude` property is not specified, the management server discovers and obtains details from all HDS storage systems managed by the discovered HiCommand Device Manager.

The IP addresses of excluded elements appear in the Discovery Data Collection lists (**Discovery > Details**). The management server does not display additional information about excluded elements in the user interface. To modify the `cimom.hds.exclude` property:

1. Select **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.

```
#cimom.hds.exclude=61038,61037
```
4. Return to the Advanced page.
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make sure the property is not commented out by removing the hash (#) symbol in front of the property. Add the serial numbers corresponding to the HDS storage systems you want to exclude from discovery. Separate additional serial numbers with a comma, as shown by the following example:

```
cimom.hds.exclude=61038,61037
```

where 61038 and 61037 are serial numbers for HDS storage systems.
7. When you are done, click **Save**.
8. Restart the service for the management server for your changes to take effect.
While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Discovering HP StorageWorks Arrays

HP CIMOM is required to discover HP StorageWorks XP Arrays, Enterprise Virtual Arrays (EVA) and Modular Smart Arrays (MSA). To discover an HP XP, EVA, or MSA storage system, you must enter the following information for the instance of the HP CIMOM used to manage the storage system. For XP storage systems, HP CIMOM is used to communicate with Command View.

- user name and password used for accessing the HP CIMOM
- IP address of the server containing the HP CIMOM

The following should be installed on a server before you discover an HP storage system:

- HP Storage Management Appliance software
- HP OpenView Storage Operations Manager
- HP StorageWorks Command View EVA, XP or MSA
- One of the following providers:
 - **XP Arrays** - HP StorageWorks SMI-S XP
 - **EVA Arrays** - HP StorageWorks SMI-S EVA
 - **MSA Arrays** - HP StorageWorks SMI-S MSA

Provisioning is supported for HP XP storage systems, but not completely for HP MSA and EVA storage systems. See [Table 42](#) on page 257 and [Table 43](#) on page 258.

To discover HP storage systems, provide the following information in HP Insight Manager:

- IP address or DNS name of the HP CIMOM you want to discover.
- User name for accessing the HP CIMOM
- Password for accessing HP CIMOM.

(XP arrays only) If you have Command View version 2.0 or later, the default password is administrator. If you have Command View earlier than version 2.0, refer to the documentation that shipped with Command View for the default password.

Discovering Engenio Storage Systems

Keep in mind the following when discovering an Engenio storage system:

- Discover all controllers on an Engenio storage system by entering the IP address of each controller.
- The management server must have the User Name field populated to discover the Engenio storage system. If your Engenio storage system does not have a user name set, you must enter something in the **User Name** field, even though the storage system has no user name.
- Discover both controllers for the Engenio storage system. Each controller has its own IP address. In Step 1 of discovery, specify all the IP addresses for all the controllers (usually two). The management server discovers these controllers as one single storage system.
- To obtain drive-related statistics, install a proxy host. Ensure the proxy host has at least one LUN rendered by each controller of the array. See the topic, "[Obtaining Disk Drive Statistics from Engenio Storage Systems](#)" on page 68 for more information.
- A license key is required for each storage system and that the key is obtained from the Web site specified on the Activation Card that shipped with your storage system.
- Engenio storage systems do not require a password for Discovery Data Collection. If you want do not want to use the management server for provisioning on Engenio storage systems, you can leave the password field blank and select the **Do Not Authenticate** option. The management server will still monitor the Engenio storage system; however, you will not be able to do provisioning tasks.

To discover Engenio storage systems, provide the following information in HP Insight Manager:

- IP address or DNS name of the controller or proxy you want to discover.
- User name for the storage system. If your Engenio storage system does not have a user name, you must enter something in the **User Name** field, even though the storage system has no user name.
- Password for the controller or proxy.

Discovering NetApp Filers

Keep in mind the following:

- SMNP must be enabled on the NetApp Filer before it can be discovered.

- If you want the management server to be able to receive events from a NetApp Filer, you must add the IP address of the management server to the NetApp configuration. The management server runs on the same computer running the management server by default.
- You must provide a privileged login, which is one of the following:
 - the root user
 - a user belonging to the “Administrators” group. This is a predefined group by NetApp.
 - a user belonging to a group that has the following roles: api-*, cli-*, login-http-admin, and at least one of the following: login-console, login-telnet, login-rsh, or login-ssh
- Administrative HTTP access to the device can be restricted through the httpd.access and httpd.admin.access options. If that is the case, then the management server needs to be registered with the device. This is done by adding the IP addresses of the management server to the httpd.admin.access option. More information related to this option is available in the NetApp documentation.

To discover a NetApp Filer, provide the following information in HP Insight Manager:

- IP address or DNS name of the NetApp Filer you want to discover.
- User name of the NetApp Filer. You must provide a privileged login.
- Password used to access the NetApp Filer.

Discovering Sun StorEdge 3510 Storage Systems

Before you can discover a Sun StorEdge 3510 storage system, you must set up a Sun StorEdge 3510 SMI-S provider and a Sun StorEdge™ Configuration Service. The provider cannot be installed on the same computer as the management server due to a port conflict.

The Sun StorEdge™ Configuration Service can be installed in one of the following locations:

- on the same computer as the Sun StorEdge 3510 SMI-S provider
- on the management server
- on a separate computer

To install the Sun StorEdge™ Configuration Service you must install the following packages:

- Sun StorEdge™ Configuration Service Console (SUNWscsu)
- Sun StorEdge™ Configuration Service Agent (SUNWscsd)
- Sun StorEdge™ Diagnostic Reporter Agent (SUNWscsa)

You must also install the following packages. Contact Sun technical support for information on how to obtain and configure these packages.

- WBEM Solutions J WBEM Server 1.0
- Sun StorEdge™ CIM/WBEM Provider SDK (SUNWagsdk package) - A readme file is installed as part of SUNWagsdk package. Follow the instructions in that readme file.
- Sun StorEdge™ 3510 SMI-S Provider (SUNW3x10a package) - A readme file is installed as part of SUNW3x10a package. Follow the instructions in that readme file.

To discover Sun StorEdge 3510 storage systems, you must discover the Sun StorEdge 3510 SMI-S provider. To discover a Sun StorEdge 3510 storage system, you must enter the following information for the instance of the Sun StorEdge 3510 SMI-S provider.

- user name and password used for the system running Sun StorEdge 3510 SMI-S provider
- IP address of the system running Sun StorEdge 3510 SMI-S provider

IMPORTANT: The management server is unable to display logical volumes configured on Sun StorEdge 3510 storage systems. Any logical volumes as well as the logical drives that comprise them will not appear in the UI. There will be no indication that this happened.

To discover Sun StorEdge 3510 storage systems, provide the following information in HP Insight Manager:

- IP address or DNS name of the system running the Sun StorEdge 3510 SMI-S provider you want to discover.
- User Name of the system running the Sun StorEdge 3510 SMI-S provider.
- Password of the Sun StorEdge 3510 SMI-S provider.

Discovering Sun StorEdge 6920 Storage Systems

To discover Sun StorEdge 6920 storage systems, provide the following information in HP Insight Manager:

- IP address or DNS name of the Sun StorEdge 6920 you want to discover.
- User name of the Sun StorEdge 6920 you want to discover.
- Password of the Sun StorEdge 6920 you want to discover.

Discovering Sun StorEdge 6130 Storage Systems

To discover Sun StorEdge 6130 storage systems, provide the following information in HP Insight Manager:

- IP address or DNS name of the controller or proxy you want to discover.
- The user name can be left blank.
- Password for the controller or proxy.

Discovering IBM Storage Systems

Before you can discover an IBM storage system, you must install:

- The IBM CIMOM, which is used to communicate with IBM storage systems. The IBM CIMOM can be installed on any host that has access to the IBM storage system. Obtain the IBM CIMOM from IBM.
- The IBM CIM Agent on a host and configured to manage one or more Enterprise Storage Server (ESS) devices. Do not install the IBM CIM Agent on the management server. Refer to the CIM Agent for the ESS - Installation and Configuration Guide for details on configuring the CIM Agent. In short, this procedure entails:

- a. Installing the software. The installation checks for the existence of the ESSCLI. If the ESSCLI is not installed, installation of the CIM Agent cannot proceed. The ESSCLI is typically pre-installed on the ESS management server that was configured by the IBM field technician.
- b. Configuring of protocol and ports used to communicate with the CIM Agent. You can change the CIM Agent port value, protocol (HTTP/HTTPS), and enable or disable the debug option.
- c. Using the `setuser` command to configure a user to access the CIM Agent. The user credentials specified here are used to access the CIMOM and are specified in the Discovery Step 1. The credentials are not necessarily the same as those used to login to the ESS Specialist management utility.
- d. Using the `setdevice` command to configure the ESS devices that are managed through the CIM Agent. The `setdevice` command requires a valid user that has the necessary privileges to access and configure the ESS storage system.
- e. Verifying that the CIM Agent is able to communicate with the ESS devices.

NOTE: Elements that were discovered through the IBM CIMOM cannot be moved to another discovery group.

To discover an IBM storage system, provide the following information in HP Insight Manager:

- In the **IP Address/DNS Name** field, enter one of the following for the system running the IBM CIMOM you want to discover.
 - `<host>` - CIM Agent has been configured to use the HTTP protocol
 - `https://<host>:5989` - CIM Agent has been configured to use the HTTPS protocol
- Provide the interop namespace as described in the online help for HP Systems Insight Manager. The namespace for an IBM storage system is usually `/root/ibm`.
- User name of the system running the IBM CIMOM.
- Password of the system running the IBM CIMOM.

Discovering IBM Tape Libraries

To discover a tape library, provide the following information in HP Systems Insight Manager:

- In the **IP Address/DNS Name** field, enter the IP address or DNS Name for the tape library.
- Provide the interop namespace as described in the online help for HP Systems Insight Manager. The namespace for an IBM tape library is usually `/root/ibm`.
- User name of the tape library.
- Password of the tape library.

Modifying the Properties of a Discovered Address

You can modify the following properties for discovering an device:

- **User name and password** - You can change the user name and password the management server uses to access a device. Whenever a user name and/or password has changed on a device the management server monitors, the management server must be made aware of the


change. For example, assume the password for a host was changed. You would need to update the management server database with the new password.

- **Discovery group** - All elements are initially placed in the Default discovery group. You can then move elements from the Default discovery group to other discovery groups. You can use discovery groups to break up Discovery Data Collection. For example, you could specify that the management server gets Discovery Data Collection for only the elements in Discovery Group 1, thus, saving you time. This feature is sometimes referred to as segmented replication because you can specify getting Discovery Data Collection for a segment of the discovered elements.

Keep in mind the following:

- You can use this window to change the user name and password stored in the management server's database. It does not change the device's user name and password.
- Discovery groups cannot be renamed or created. You must use the existing discovery groups.
- You can also use the **Move to Discovery Group** button to move multiple elements to another discovery group. See "[Moving Elements to Another Discovery Group](#)" on page 45 for more information.

To change the discovery properties of an element:

1. Click **Options > Storage Essentials > Discovery > Run Discovery Data Collection**.
2. Click the  button corresponding with the element you want to modify.
3. To change the user name, type the new user name in the **User Name** field.
4. To move an element to another discovery group, select its new discovery group from the **Discovery Group** drop-down menu.
5. To change the password:
 - a. Click **Change Password**.
 - b. Type the new password in the **New Password** field.
 - c. Type the password again in the **Verify Password** field.
 - d. Click **OK** in the Change Password window.
6. Click **OK** in the Edit Discovered Element window.

Deleting Elements from the Management Server

When you delete an element, all of its information is removed from the management server. This includes asset information, zoning, events, statistics, and fabrics assigned to switches.


To completely delete an element from the management server you must remove the elements, such as a switch or proxy that were used to discover the element. If you do not delete all switches and proxies that were used to discover the element, the element may reappear the next time you Discovery Data Collection.

For example, assume you want to delete Switch_A. Switch_B and Switch_C were used to discover Switch_A. If you delete only Switch_B and Switch_A, Switch_A will most likely reappear when you Discovery Data Collection because it is still accessible by Switch_C.

Deleting an Element Using System Manager or Chargeback Manager

To delete an element using System Manager or Chargeback Manager:

1. Do one of the following:
 - **In System Manager** - Right-click an element and select **Delete Element** from the drop-down menu. Right-click an element and select **Delete Element** from the drop-down menu.

If you are blocking pop-ups and you use the right-click menu to delete an element from System Manager, the Delete window is blocked and you are unable to delete the element. You must disable the popup blocker before you can delete the element.
 - **In Chargeback Manager** - Click the  button for the element you want to delete.
2. If the element has multiple access points, you are asked which want to delete. Do one of the following:
 - **Delete the element and its access points.** This option lists not only the switch you want to delete, but also the other elements that use the same switches and proxies as the element you want to delete. For example, assume you want to delete Switch_A. Switch_B was used to discover Switch_A. Let's assume Switch_B is also the only path to Switch_D. If you delete Switch_B, you will no longer have access to Switch_D. This option would list Switch_D as one of the other elements that need to be deleted.

An access point is the intersection of the IP address and the provider that discovered the IP address. A provider is software that is used to gather information about an element.
 - **Delete the element.** The element may reappear the next time you obtain element details. This is because not all switches and proxies connected to the element have not been removed. For example, assume you want to delete Switch_A. Switch_B is connected to Switch_A. If you do not delete Switch_B, the next time you obtain element details Switch_B will most likely find Switch_A again.
3. Click **OK**.

Step 3 - Discovery Data Collection

This section describes the following:

- ["Discovery Data Collection" on page 39](#)
- ["Stopping the Gathering of Details" on page 40](#)
- ["Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh" on page 41](#)
- ["Excluding HDS Storage Systems from Force Device Manager Refresh" on page 42](#)

Discovery Data Collection

Discovery Data Collection is required to obtain detailed information from discovered elements. Discovery Data Collection must be performed before you can do provisioning and/or obtain provisioning information, such as data about zone sets and LUN numbers.

Keep in mind the following:

- Discovery Data Collection takes some time. You might want to perform this process when the network and the managed elements are not busy.
- If you have problems obtaining information from Connectrix and McDATA switches during Discovery Data Collection, see the topic, ["Step 2 - Discover Storage Systems, Filers and Tape Libraries"](#) on page 26.
- You can use discovery groups to break up getting Discovery Data Collection. For example, instead of Discovery Data Collection for all of the elements, you could specify that the management server Discovery Data Collection for only the elements in Discovery Group 1, thus, saving you time. You add an element to a discovery group by modifying the properties used to discover the element. See ["Moving Elements to Another Discovery Group"](#) on page 45 for information on how to move one or more multiple elements to a discovery group. You can also move an element to another discovery group when you modify its discovery properties. See ["Modifying the Properties of a Discovered Address"](#) on page 37.
- When an element in a given discovery group is updated, its dependent elements are also updated. For example, assume Host_A is the only element in Discovery Group 1. Host_A is connected through a switch and storage system. When you Discovery Data Collection for Discovery Group 1, you also obtain details from the switch and storage system.
- You can quarantine elements to exclude them from Discovery Data Collection. See ["Placing an Element in Quarantine"](#) on page 46 for more information. Let us assume you want to discover all the elements in a discovery group, except for one. Perhaps the element you want to quarantine is being taken off the network for maintenance. You can use the quarantine feature to exclude one or more elements from discovery.
- If you want to receive status reports about Discovery Data Collection, see ["Configuring E-mail Notification for Discovery Data Collection"](#) on page 511 for information about how to configure this option.
- If the management server unable to obtain information from a UNIX host during Discovery Data Collection as a result of a CIM Extension hanging, the management server places the access point where the CIM Extension is located in quarantine. The management server then moves onto getting details for the next element in the Discovery Data Collection table. These UNIX hosts appear as missing until they are removed from quarantine. See ["Removing an Element from Quarantine"](#) on page 46 for information on how to remove an element from quarantine.

To obtain details about the devices on the network:

1. Click **Options > Storage Essentials > Discovery > Run Discovery Data Collection**.
2. Select **Include infrastructure details**, which gathers information about SAN details.
3. The management server obtains most of the information from HDS and EMC Symmetrix storage systems from their device managers. Select **Force Device Manager Refresh** if you want the management server to tell the device managers for HDS and EMC storage systems to obtain the latest information. See the following topics for more information: ["Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh"](#) on page 41 and ["Excluding HDS Storage Systems from Force Device Manager Refresh"](#) on page 42.

NOTE: If you plan to have File SRM scan a host, make sure you have 220 to 230 MB for each set of 1 million files.

4. Select the discovery group from which you want to obtain Discovery Data Collection. If you are obtaining Discovery Data Collection for the first time, make sure **All Discovery Groups** is selected.
5. Click the **Get Details** button.

While getting element details, the software changes its status light from green to red. You can view the progress of gathering details by clicking **Discovery > View Logs**.

When the software completes getting all elements details, it displays "GETTING ALL DETAILS COMPLETED" on the **View Logs** page.
6. See the "Adding a Discovery Schedule" on page 100 for information about automating the gathering of all element details.
7. To add more IP addresses, IP Ranges or application information for discovery before completing the following step, click the **Discovery > Setup** link displayed below the logs screen.

Once you add more elements to be discovered, obtain Discovery Data Collection.

Stopping the Gathering of Details

Obtaining details takes some time. If the network and managed elements are busy, you might need to stop the gathering of details and reschedule it for another time.

IMPORTANT: If you stop the gathering of details, you should reschedule it. This type of collection obtains detailed information of devices in the network.

To stop the gathering of details:

1. Click **Options > Storage Essentials > Discovery > Run Discovery Data Collection**.
2. On the **View Logs** tab, click the "Click here" portion of the following message:

Click here if you wish to stop getting details.
3. When you are asked if you are sure you want to stop Discovery Data Collection, click **OK**.

The management server stops gathering details.
4. Schedule a time to resume getting details.

Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh

The management server obtains most of its information about Symmetrix storage systems from the EMC Solutions Enabler (proxy server) it discovered. If the EMC Solutions Enabler does not have the latest information, the management server also displays the outdated information.

To make the management server aware of any changes, make sure the Solutions Enabler it discovered has the latest information. This can be done by forcing the Solutions Enabler to refresh its data. The management server is then made aware of these changes.

When the **Force Device Manager Refresh** option is selected, the management server refreshes discovered EMC Solutions Enabler (proxy server), unless specified. If you do not want an EMC Solutions Enabler to be refreshed, you must assign the Symmetrix storage systems that use the Solutions Enabler to the `cimom.emc.skipRefresh` property, as described in the steps in this section.

To exclude EMC Symmetrix storage systems from a forced refresh:

1. Select **Options >Storage Essentials > Manage Product Health > Advanced**.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
`#cimom.emc.skipRefresh=000183500570,000183500575`
4. Return to the Advanced page (**Options >Storage Essentials > Manage Product Health > Advanced**).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make sure the property is not commented out by removing the hash (#) symbol in front of the property. Add the serial numbers corresponding to the Symmetrix storage systems you want the refresh to skip. Separate additional serial numbers with a comma, as shown by the following example:
`cimom.emc.skipRefresh=000183500570,000183500575`
where 000183500570 and 000183500575 are serial numbers for Symmetrix storage systems. One of the ways to find the serial number is to double-click the storage system in System Manager. Then, click the **Properties** tab.
7. When you are done, click **Save**.
8. Restart the service for the management server for your changes to take effect:
 - a. Go to the Services window on the management server.
 - b. Right-click **AppStorManager**.
 - c. Select **Restart** from the drop-down menu.While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.
9. To perform the forced refresh, select the **Force Device Manager Refresh** option on the Discovery Data Collection page (**Discovery > Details**).
10. Click **Get Details**.

Excluding HDS Storage Systems from Force Device Manager Refresh

The management server obtains most of its information about the HDS storage systems from the HiCommand Device Manager (proxy server) it discovered. If HiCommand Device Manager, does not have the latest information, the management server also displays the outdated information.

To make the management server aware of any changes, make sure the HiCommand Device Manager it discovered has the latest information. This can be done by forcing the HiCommand Device Manager to refresh its data. The management server is then made aware of these changes.

When the **Force Device Manager Refresh** option is selected, the management server refreshes discovered HiCommand Device Manager (proxy server), unless specified. If you do not want a HiCommand Device Manager to be refreshed, you must assign the HDS storage systems that use HiCommand Device Manager to the `cimom.HdsSkipRefresh` property, as described in the steps in this section.

IMPORTANT: Before performing any provisioning operations, you should perform a forced refresh.

To exclude HDS storage systems from a forced refresh:

1. Select **Options >Storage Essentials > Manage Product Health > Advanced**.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.

```
# cimom.HdsSkipRefresh=61038,61037
```
4. Return to the Advanced page (**Options >Storage Essentials > Manage Product Health > Advanced**).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make sure the property is not commented out by removing the hash (#) symbol in front of the property. Add the serial numbers corresponding to the HDS storage systems you want the refresh to skip. Separate additional serial numbers with a comma, as shown by the following example:

```
cimom.HdsSkipRefresh=61038,61037
```

where 61038 and 61037 are serial numbers for HDS storage systems. One of the ways to find the serial number is to double-click the storage system in System Manager. Then, click the **Properties** tab.
7. When you are done, click **Save**.
8. Restart the service for the management server for your changes to take effect:
 - a. Go to the Services window on the management server.
 - b. Right-click **AppStorManager**.
 - c. Select **Restart** from the drop-down menu.

While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

9. To perform the forced refresh, select the **Force Device Manager Refresh** option on the Discovery Data Collection page (**Discovery > Details**).

10. Click **Get Details**.

Managing McDATA and EMC Connectrix Switches

This section describes the following:

- “[About Managing McDATA and EMC Connectrix Switches](#)” on page 43
- “[Adding McDATA and EMC Connectrix Switches](#)” on page 43

About Managing McDATA and EMC Connectrix Switches

Whenever you add, McDATA or EMC Connectrix switches in an already discovered service processor, you must make the management server aware of those changes. After you add these switches to the service processor, you must perform “Discovery Data Collection” in the management server. The management server obtains information about the new switches from the service processor. See the topic, “[Adding McDATA and EMC Connectrix Switches](#)” on page 43 for more information about adding switches.

Adding McDATA and EMC Connectrix Switches

After you add switches to an existing service processor, you must perform Discovery Data Collection, as described in the following steps. If you are adding switches to a service processor that has not been discovered yet, see the topic, “[Discovering McDATA and EMC Connectrix Switches](#)” on page 18.

IMPORTANT: Obtaining details takes some time. You might want to perform this process when the network and the managed elements are not busy.

To Discovery Data Collection:

1. Click **Options > Storage Essentials > Discovery > Run Discovery Data Collection**.
2. Click the **Get Details** button.

While getting element details, the software changes its status light from green to red.

Assigning a File Extension in Netscape 7

Netscape 7 automatically assigns unknown files an HTML extension. To make Netscape 7 recognize the type of file, you must assign a file extension.

To assign a MIME type:

1. Click the download file link or button in the software.
2. Click the **Advanced** button in the lower-left corner.
3. In the **Description of type** field, delete the existing text and type a description of the file.

4. In the **File extension** field, delete the existing text and type the file extension.
5. Click **OK**.

The next time Netscape 7 sees the associated MIME type, it will assign the extension you typed in the **File Extension** field.

For example, in the following figure, the zip extension was assigned to a MIME type of application/unknown. The next time Netscape sees that MIME type, it will automatically assign the zip extension to the file.

6. Click **OK**.

Updating the Database with Element Changes

After you have initially discovered the elements, information about them might change. To update database with these changes, perform the steps described in this section.

Keep in mind the following:

- If you change the password of a host after you discover it, you must change the password for the host in the discovery list. Then, you must stop and restart the CIM Extension running on that host.
- If you are adding McDATA or Connectrix switches, you must perform different steps. See the topic, "[Adding McDATA and EMC Connectrix Switches](#)" on page 43.

To update the database:

1. Click **Options > Storage Essentials > Discovery > Run Discovery Data Collection**.
2. Make sure the File Server SRM option is selected.
3. Select **Include infrastructure details**, which gathers information about SAN details.
4. The management server obtains most of the information from HDS and EMC Symmetrix storage systems from their device managers. Select **Force Device Manager Refresh** if you want the management server to tell the device managers for HDS and EMC storage systems to obtain the latest information. See the following topics for more information: "[Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh](#)" on page 41 and "[Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh](#)" on page 41.
5. Click the **Get Details** button on the Discovery Data Collection page.
6. View the status of the gathering of element details by looking in the **View Logs** tab. See the topic, "[Viewing Log Messages](#)" on page 25 for more information about the messages viewed in this tab.
7. Verify the topology is displayed correctly by accessing System Manager. To access System Manager, click the **System Manager** button in the left pane.

Filtering Discovery Groups

You can determine which discovery groups are displayed on the Discovery Data Collection (**Options > Storage Essentials > Discovery > Run Discovery Data Collection**) page by modify the discovery filter, as described in the following steps:

1. Access the Discovery Data Collection page (**Options > Storage Essentials > Discovery > Run Discovery Data Collection**).

2. Click the **Custom** button.
3. Select the discovery groups you want to include in Discovery Data Collection. Deselect the discovery groups you do not want to be included in Discovery Data Collection.
4. Click **OK**.

Elements in the selected discovery groups are selected on the Discovery Data Collection page. The management server obtains information from the selected elements during Discovery Data Collection. To learn how to add an element to a different discovery group, see "[Modifying the Properties of a Discovered Address](#)" on page 37.

Moving Elements to Another Discovery Group

All elements are initially placed in the Default discovery group. You can then move elements from the Default discovery group to other discovery groups. You can use discovery groups to break up getting Discovery Data Collection. For example, you could specify that the management server gets Discovery Data Collection for only the elements in Discovery Group 1, thus, saving you time. This feature is sometimes referred to as segmented replication because you can specify getting Discovery Data Collection for a segment of the discovered elements.

Keep in mind the following:

- Discovery groups cannot be renamed or created. You must use the existing discovery groups.
- You can also use move an element to another discovery group when you modify its discovery properties. See "[Modifying the Properties of a Discovered Address](#)" on page 37 for more information.

To move an element to another discovery group:

1. Select the check boxes for the elements you want to move in the Discovery Data Collection page.
2. Click the **Move to Discovery Group** button.
3. In the Select Discovery Group window, select the new discovery group for the selected elements.
4. Click **OK**.

The elements are moved to the new discovery group.


Placing an Element in Quarantine

When you click the **Get Details** button on the Discovery Data Collection page, the management server automatically obtains details for the elements in the selected discovery group. Let us assume you want to discover all the elements in a discovery group, except for one. Perhaps the element you want to quarantine is being taken off the network for maintenance. You can use the quarantine feature to exclude one or more elements from discovery.

NOTE: After you perform Discovery Data Collection for the discovery group containing the quarantined elements, the quarantined elements appear as missing throughout the product. The management server marks the quarantined elements as missing because it cannot obtain details from the quarantined element.

To quarantine an element:

1. Select the check boxes for the elements you want to quarantine on the Discovery Data Collection page.
2. Click the **Set Quarantine** button.
3. When you are asked if you want to quarantine the selected elements, click **OK**.


The elements you quarantine appear with a flag () in the Quarantined column on the Discovery Data Collection page.

The elements are excluded from discovery until you clear them from quarantine.

Removing an Element from Quarantine

To remove an element from quarantine:

1. Select the check boxes for the elements you want to remove from quarantine on the Discovery Data Collection page.

Quarantined elements appear with a flag () in the Quarantined column on the Discovery Data Collection page.

2. Click the **Clear Quarantine** button.
3. When you are asked if you want to remove the selected elements from quarantine, click **OK**.

The next time you perform Discovery Data Collection for the element, the management server gathers data from the element.

3 Discovering Applications and Hosts

This chapter describes the following:

- ["Step 1 - Discovering Your Hosts"](#) on page 49
- ["Step 2 - Setting Up Discovery for Applications"](#) on page 51
- ["Step 3 - Discovering Applications"](#) on page 64
- ["Changing the Oracle TNS Listener Port"](#) on page 65
- ["Adding/Modifying Microsoft Exchange Domain Controller Access"](#) on page 66
- ["Deleting a Microsoft Exchange Domain Controller"](#) on page 67
- ["Changing the Password for the Managed Database Account"](#) on page 67
- ["Obtaining Disk Drive Statistics from Engenio Storage Systems"](#) on page 68
- ["Assigning a File Extension in Netscape 7"](#) on page 68

Step 1 - Discovering Your Hosts

Before you can discover your applications, you must discover their hosts. You discover hosts in the same way you discovered your switches and storage systems. You provide the host's IP address, user name and password. The user name and password must have administrative privileges. Unlike switches and storage systems, you must have installed CIM Extension on the host if you want to obtain detailed information about the host.

Keep in mind the following:

- If you change the password of a host after you discover it, you must change the password for the host in the discovery list. Then, you must stop and restart the CIM Extension running on that host.
- If your license lets you discover UNIX and/or Linux hosts, the **Test** button for discovery reports SUCCESS from any UNIX and/or Linux hosts on which the management server can detect a CIM Extension. The CIM Extension must be running. The management server reports "SUCCESS" even if your license restricts you from discovering certain types of hosts. For example, assume your license lets you discover Solaris hosts but not AIX hosts. If you click the **Test** button, the management server reports "SUCCESS" for the AIX hosts. You will not be able to discover the AIX hosts. The IP address is not discoverable, because of the license limitation.
- You should have already installed a CIM Extension on the host you want to discover.
- If you want to receive status reports about Discovery Data Collection, see ["Configuring E-mail Notification for Discovery Data Collection"](#) on page 511 for information about how to configure this option.
- Depending on your license, you may not be able to access File System Viewer and/or monitor certain applications may not be available. See the List of Features to determine if you have access to File System Viewer and/or are able to monitor the other applications. The List of Features is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials). To learn more about File System Viewer, refer to the File Servers Guide, which is also available from the Documentation Center.

- If you are unable to discover a UNIX host because of DNS or routing issues, see [“Unable to Discover a UNIX Host Because of DNS or Routing Issues”](#) on page 520.

Discovery of hosts consists of two steps:

- Detecting the hosts by using HP Systems Insight Manager. See [“Step A - Set Up Discovery for Hosts”](#) on page 50.
- Obtaining details about those hosts by using the Discovery Data Collection feature in Storage Essentials. See [“Step B - Obtain Details”](#) on page 50.

Step A - Set Up Discovery for Hosts

IMPORTANT: Dynamic disk support on Windows 2000 hosts is optional. If you want this feature enabled for Windows 2000 hosts, you can download LDMDump, as described in [“Enabling Dynamic Disk Detection for Windows 2000 Hosts”](#) on page 50.

Use HP Systems Insight Manager (SIM) to discover your hosts. Refer to your documentation for HP SIM for more information.

Enabling Dynamic Disk Detection for Windows 2000 Hosts

Dynamic disk support for Windows 2000 hosts is optional. If you want this feature enabled for Windows 2000 hosts, perform the following steps for the management server to obtain information about dynamic disks on that host:

1. Download LDMDump from <http://www.sysinternals.com/Utilities/LdmDump.html> to the Windows 2000 host.
2. Unzip the LdmDump.zip file into the Windows\System32 folder on the Windows 2000 host.
3. Discover the Windows 2000 host.
4. If the management server detects LDMDump, the following is displayed in the Log Messages window.

Dynamic disk volumes supported via the utility 'ldmdump': Logical Disk Manager Configuration Dump v1.03

NOTE: The version number displayed may vary from the version you downloaded.

Step B - Obtain Details

Discovery Data Collection must be performed before you can do provisioning and/or obtain provisioning information, such as data about zone sets and LUN numbers.

Keep in mind the following:

- Obtaining details takes some time. You might want to perform this process when the network and the managed elements are not busy.
- If you want to enable File System Viewer for a host, make sure the File SRM option is selected.
- If Discovery Data Collection includes an AIX host, the system log displays three SCSI errors (2 FSCSI error and 1 FCS error) per IBM adapter port. You can ignore these errors.

- You can quarantine elements to exclude them from Discovery Data Collection. See [“Placing an Element in Quarantine”](#) on page 46 for more information. Let us assume you want to discover all the elements in a discovery group, except for one. Perhaps the element you want to quarantine is being taken off the network for maintenance. You can use the quarantine feature to exclude one or more elements from discovery.
- When an element in a given discovery group is updated, its dependent elements are also updated. For example, assume Host_A is the only element in Discovery Group 1. Host_A is connected through a switch and storage system. When you Discovery Data Collection for Discovery Group 1, you also obtain details from the switch and storage system.
- If the management server unable to obtain information from a UNIX host during Discovery Data Collection as a result of a CIM Extension hanging, the management server places the access point where the CIM Extension is located in quarantine. The management server then moves onto getting details for the next element in the Discovery Data Collection table. These UNIX hosts appear as missing until they are removed from quarantine. See [“Removing an Element from Quarantine”](#) on page 47 for information on how to remove an element from quarantine.

To obtain details:

1. Select **Options > Storage Essentials > Discovery > Run Discovery Data Collection**.
2. Verify the **File SRM** option is selected. The File SRM option appears for hosts that have the CIM Extension and an operating system that supports File System Viewer.
3. Click the **Get Details** button.

For additional information, see [“Updating the Database with Element Changes”](#) on page 45 for information on how to automate the gathering of all element details. If you run into problems with discovery, see [“Troubleshooting”](#) on page 501.

Step 2 - Setting Up Discovery for Applications

Keep in mind the following when discovering applications:

- Make a list of the applications you want to monitor. Configure your applications first as described in this section and then run discovery.
- You should have already installed a CIM Extension on the hosts that have the applications you want to discover. After you installed the CIM Extension, you should have already discovered the host. See [“Step 1 - Discovering Your Hosts”](#) on page 49.

You can configure the management server to monitor hosts and applications, such as Oracle, Microsoft Exchange server, and Sybase Adaptive Server Enterprise, in addition to Microsoft SQL servers and file servers. If you want to obtain detailed information about the host and its applications, you must install a CIM Extension on the host, as described in the previous chapters.

The following is an overview of what you need to do. It is assumed you have already discovered the hosts running your applications. See [“Step 1 - Discovering Your Hosts”](#) on page 49.

Then, set up the configurations for your applications on the management server. Some applications may require you to provide additional discovery information about the application. Finally, perform discovery and then obtain Discovery Data Collection. Obtaining details takes some time. Perform this step when the network is not busy. More details about the steps mentioned above are provided later.

See the following topics for more information:

- ["Monitoring Oracle" on page 52](#)
- ["Monitoring Microsoft SQL Server" on page 57](#)
- ["Monitoring Sybase Adaptive Server Enterprise" on page 59](#)
- ["Monitoring Microsoft Exchange" on page 62](#)

Monitoring Oracle

To monitor and manage Oracle, you must do the following:

- ["Step A - Create the APPIQ_USER Account for Oracle" on page 52](#)
- ["Step B - Provide the TNS Listener Port" on page 54](#)
- ["Step C - Set up Discovery for Oracle 10g" on page 55](#)

After you complete these steps, you must discover Oracle and obtain Discovery Data Collection. See ["Step 3 - Discovering Applications" on page 64](#).

IMPORTANT: Before you begin these steps, make sure you purchased the module that lets you monitor Oracle. Contact your customer support if you are unsure if you purchased this module.

Step A - Create the APPIQ_USER Account for Oracle

The management server accesses Oracle through the APPIQ_USER account. This account is created when you run the `CreateOracleAct.bat` script on Microsoft Windows or `CreateOracleAct.sh` on UNIX on the computer running the Oracle database you want to monitor. The account has create session and select dictionary privileges to be used with the management server.

Keep in mind the following:

- The `CreateOracleAct.bat` script must run under SYS user.
- Create APPIQ_USER account on Oracle Database you want to monitor, not on the management server.
- You should have already installed the database for the management server.
- Verify that the instance TNS (Transparent Name Substrate) listener is running so that the management server can find the Oracle installation and its instances. For example on Microsoft Windows 2000, you can determine if the instance TNS listener is running by looking in the Services window for OracleOraHome92TNSListener. The name of the TNS listener might vary according to your version of Oracle. Refer to the Oracle documentation for information about verifying if the instance TNS listener is running. You can also verify the listener is running by entering the following at the command prompt: `lsnrctl status`. If the listener is not running you can start it by typing `lsnrctl start` on command line.
- Make sure you have all the necessary information before you begin the installation. Read through the following steps before you begin.

To create the Oracle user for management server:

1. Do one of the following:

- **To run the script on IBM AIX, SGI IRIX, or Sun Solaris**, log into an account that has administrative privileges, mount the CD-ROM (if not auto-mounted), and go to the `/DBIQ/Oracle/unix` directory by typing the following:

```
# cd /cdrom/cdrom0/DBIQ/Oracle/unix
```

where `/cdrom/cdrom0` is the name of the CD-ROM drive
- **To run the script on Microsoft Windows**, go to the `DBIQ\Oracle\win` directory on the CD-ROM.

IMPORTANT: You must complete the following steps.

2. Verify you have the password to the SYS user account.

You are prompted for the password for this user account when you run the script.

3. Run the `CreateOracleAct.bat` script on Microsoft Windows or `CreateOracleAct.sh` script on the UNIX operating system on the computer with the Oracle database.

The script creates a user with create session and select dictionary privilege on a managed Oracle instance.

NOTE: You can use a remote Oracle client to run this script.

4. Specify the Oracle instance name, which must be visible to the client, as the first input when running the script. The script prompts you for the name of the Oracle instance on which to create user for Oracle management packages and the password of the SYS account.

You are asked to specify the default and temporary tablespaces for APPIQ_USER during the installation. You can enter users as default and temp as temporary if these tablespaces exist in the Oracle Instance.

5. Repeat the previous step for each Oracle instance you want to manage.

This script does the following in order:

- Creates the APPIQ_USER account.
- Grant create session and select on dictionary tables privileges to APPIQ_USER enabling management server to view statistics for the Oracle instances.

Removing the APPIQ_USER Account for Oracle

If you no longer want the management server to monitor an Oracle instance, you can remove the APPIQ_USER account for that Oracle instance by running the `UninstallOracleAct.bat` script on Windows or `UninstallOracleAct.sh` script on the UNIX platform.

Keep in mind the following:

- Before you remove the APPIQ_USER account for an Oracle instance, make sure no processes are running APPIQ_USER for that Oracle instance. The management server uses APPIQ_USER to obtain information about the Oracle database. For example, a process would be using APPIQ_USER if someone was using Performance Manager to view monitoring statistics about

that Oracle instance. After you removed the APPIQ_USER account for Oracle, discover and perform Discovery Data Collection for the host if you want to continue monitoring it.

- If you receive a message about not being able to drop a user that is currently connected while you are removing the APPIQ_USER account for Oracle, re-run the script for removing APPIQ_USER.

To remove the APPIQ_USER account for that Oracle instance:

1. If you plan to remove the management software for Oracle from the Solaris host, do the following:
 - a. Log into an account that has administrative privileges.
 - b. Mount the CD-ROM (if not auto-mounted)
 - c. Go to the `/DBIQ/Oracle/unix` directory by typing the following:

```
# cd /cdrom/cdrom0/DBIQ/Oracle/unix
```

where `/cdrom/cdrom0` is the name of the CD-ROM drive
2. If you plan to remove the management software for Oracle from a computer running Windows, go to the `\DBIQ\Oracle\win` directory on the CD-ROM.
3. Verify you have the password to the SYS user account.

You are prompted for the password for this user account when you run the script.
4. Run the `UninstallOracleAct.bat` for Windows or `UninstallOracleAct.sh` script for UNIX platform on the computer with the Oracle database.
5. This script removes the management software for the specified Oracle instance.

NOTE: You can use a remote Oracle client to run this script.

6. When you are asked for the Oracle instance name, enter the name of the Oracle instance you do not want the management server to monitor. The name must be visible to the client.
7. Provide the password for the SYS user account.


The APPIQ_USER account for the specified Oracle instance is removed. The management server can no longer monitor that Oracle instance.

Step B - Provide the TNS Listener Port

If your Oracle instances use a different TNS Listener Port than 1521, change the port as described in the following steps:

1. Select **Options > Protocol Settings > Storage Essentials > Global Application Discovery Settings**.

The TNS Listener Port setting applies to all Oracle instances you monitor.

2. To assign a new port, click the **Create** button for the Oracle Information table.
3. Type the new port number and click **OK**.
4. If necessary, click the  button to remove the old port number.

IMPORTANT: Monitoring Oracle 10g and Oracle clusters require an additional step. If you are not monitoring Oracle 10g and Oracle clusters, see ["Step 3 - Discovering Applications"](#) on page 64.

Step C - Set up Discovery for Oracle 10g

NOTE: If you are discovering an Oracle cluster, see ["Discovering Oracle Clusters"](#) on page 55.

To monitor Oracle 10g, provide additional information as described in the following steps:

1. Select **Options > Protocol Settings > Storage Essentials > System Application Discovery Settings**.

To select a target, you must have at least one element designated as a server, workstation or desktop. If you see the message, "No Targets Currently Selected," change your element from unknown to either a server, workstation or desktop. Refer to the documentation for HP Systems Insight Manager.

2. Select a target, and then, click **Run Now**.

3. Click the **Create** button for the Database Information table.

4. In the **Host IP/DNS Name** field, type the IP address or DNS name of the host running Oracle. The **Management IP/DNS Name** field is optional.

5. In the **Server Name** field, type the Oracle System Identifier (SID) of the Oracle database you want to monitor.

6. In the **Port Number** field, type the monitored port.

If you are not sure of the monitored port, check the listener.ora file of the monitored database application. You can find the listener.ora file in the following directory on the host of the monitored database. Do not look for the listener.ora file on the management server for this information.

%ORA_HOME%\network\admin\listener.ora

The port can be found in the following code:

```
LISTENER =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS_LIST =
        (ADDRESS = (PROTOCOL = TCP) (HOST = localhost) (PORT = 1521))
        (ADDRESS = (PROTOCOL = IPC) (KEY = EXTPROC0))
      )
    )
  )
```

7. Select **ORACLE** from the Database Type menu.

8. Click **OK**.

Discovering Oracle Clusters

Perform the following steps for each node in the cluster:

1. Install the CIM Extension on each node in the cluster. See the Installation Guide for information on how to install the CIM Extensions. See ["Roadmap for Installation and Initial Configurations"](#) on page 2 for information about the CIM Extensions available.
2. Create the appiq_user account on each node in the cluster. See ["Step A - Create the APPIQ_USER Account for Oracle"](#) on page 52.
3. Discover the host for the first node.
4. Discover first Oracle node as follows:

- a. Select **Options > Protocol Settings > Storage Essentials > System Application Discovery Settings**.

To select a target, you must have at least one element designated as a server, workstation or desktop. If you see the message, "No Targets Currently Selected," change your element from unknown to either a server, workstation or desktop. Refer to the documentation for HP Systems Insight Manager.

- b. Select a target, and then, click **Run Now**.
- c. Click the **Create** button for the Database Information table.
- d. In the **Host IP/DNS Name** field, type the IP address or DNS name of the host running Oracle.

In the **Management IP/DNS Name** field, type the IP address the listener is listening on for the Oracle instance. The IP address can be a virtual IP or a host IP. You can find the IP address in the listener.ora file for the monitored database. You can find the listener.ora file in the following directory on the host of the monitored database. Do not look for the listener.ora file on the management server for this information.

```
%ORA_HOME%\network\admin\listener.ora
```

- e. In the **Server Name** field, type the Oracle System Identifier (SID) of the Oracle database you want to monitor.
- f. In the **Port Number** field, type the monitored port.

If you are not sure of the monitored port, check the listener.ora file of the monitored database application. You can find the listener.ora file in the following directory on the host of the monitored database. Do not look for the listener.ora file on the management server for this information.

```
%ORA_HOME%\network\admin\listener.ora
```

The port can be found in the following code:

```
LISTENER =  
  (DESCRIPTION_LIST =  
    (DESCRIPTION =  
      (ADDRESS_LIST =  
        (ADDRESS = (PROTOCOL = TCP) (HOST = localhost) (PORT = 1521))  
        (ADDRESS = (PROTOCOL = IPC) (KEY = EXTPROC0))  
      )  
    )  
  )  
)
```

g. Select **ORACLE** from the Database Type menu.

h. Click **OK**.

5. Repeat Steps 4 and 5 for each node in the cluster.

Monitoring Microsoft SQL Server

To manage and monitor Microsoft SQL Servers, you must do the following:

- ["Step A - Create the APPIQ_USER for the SQL Server"](#) on page 57
- ["Step B - Provide the Microsoft SQL Server Name and Port Number"](#) on page 58

IMPORTANT: Make sure the Microsoft SQL server database is in "Mixed Mode authentication." To switch to mixed mode authentication, see ["Switching to Mixed Mode Authentication"](#) on page 57.

Switching to Mixed Mode Authentication

Microsoft SQL Server must be running in Mixed Mode Authentication. You can switch to Mixed Mode Authentication as follows:

1. Open SQL Server Enterprise Manager (Start menu > **Programs** > **Microsoft SQL Server** > **Enterprise Manager**).
2. Expand the tree-control until you can see your server.
3. Right-click the server name.
The SQL Server Properties (Configure) window appears.
4. Click the **Security** tab.
5. For "Authentication", select **SQL Server and Windows**.

Step A - Create the APPIQ_USER for the SQL Server

The management server accesses SQL Server through the APPIQ_USER account. This account is created when you run the `CreateSQLServerAct.bat` script on Microsoft Windows on the computer running the SQL Server database you want to monitor. The account has create session and select dictionary privileges to be used with the management server.

Keep in mind the following:

- The script must run under SA user.

- Obtain the SQL Server name before you run the script
- Create APPIQ_USER account on SQL Server database you want to monitor.
- You should have already installed the database for the management server.
- Make sure you have all the necessary information before you begin the installation. Read through the following steps before you begin.

To create the APPIQ_USER account for SQL Server:

1. To run the script on Microsoft Windows, go to the `DBIQ\sqlserver\win` directory on the CD-ROM.

IMPORTANT: You must complete the following steps.

2. Verify you have the password to the SA user account.
You are prompted for the password for this user account when you run the script.
3. Run the `CreateSQLServerAct.bat` script on Microsoft Windows on the computer with the SQL Server database.
The script creates a user with login to master and select privilege on data dictionary tables on a managed SQL Server instance.

NOTE: You can use a remote SQL Server `isql` to run this script.

4. Type the SQL Server instance name, which must be visible to the client, as the first input when running the script. The script prompts you for the name of the SQL Server on which to create user for SQL Server management packages and the password of the SA account.
5. Repeat the previous step for each SQL Server you want to manage.
This script does the following in order:
 - Creates the APPIQ_USER account.
 - Grant create session and select on dictionary tables privileges to APPIQ_USER enabling management server to view statistics for the SQL Server.

Removing the APPIQ_USER Account for SQL Server

IMPORTANT: Before you remove the APPIQ_USER account for the SQL Server databases on a host, make sure no processes are running APPIQ_USER for that SQL Server database. The management server uses APPIQ_USER to obtain information about a SQL Server database.

To remove the APPIQ_USER account for the SQL Server databases on a host:

1. To run the script on Microsoft Windows, go to the `DBIQ\sqlserver\win` directory on the CD-ROM.

IMPORTANT: You must complete the following steps.

2. Verify you have the password to the SA user account.
You are prompted for the password for this user account when you run the script.
3. Run the `UninstallSQLServerAct.bat` script on Microsoft Windows on the computer with the SQL Server database.
4. Type the name of the SQL Server server.
5. Type the password for the SA account.
The account for APPIQ_USER is removed. The management server can no longer monitor the SQL Server databases on this host.

Step B - Provide the Microsoft SQL Server Name and Port Number

You must provide the server name for the Microsoft SQL server and port number for managing a SQL database in the following steps:

To add information for discovering a SQL server:

1. Select **Options > Protocol Settings > Storage Essentials > System Application Discovery Settings**.
To select a target, you must have at least one element designated as a server, workstation or desktop. If you see the message, "No Targets Currently Selected," change your element from unknown to either a server, workstation or desktop. Refer to the documentation for HP Systems Insight Manager.
2. Select a target, and then, click **Run Now**.
3. Click the **Create** button for the Database Information table.
4. In the **Host IP/DNS Name** field, type the IP address or DNS name of the host running Sybase.
5. You can leave the **Management IP/DNS Name** field blank. This field is for Oracle clusters. When you leave the **Management IP/DNS Name** field blank the management server automatically lists the DNS name or IP address of the host under the **Host IP/DNS Name** column and **Management IP/DNS Name** column.
6. In the **Server Name** field, type the SQL database you want to monitor.
7. In the **Port Number** field, type the port that SQL is using. If you do not enter a port number, the management server assumes you are using port 1433 (default).
8. Select **SQLSERVER** from the Database Type menu.
9. Click **OK**.


IMPORTANT: Perform Discovery Data Collection for your inputs to take effect. See "[Step 3 - Discovering Applications](#)" on page 64.

Deleting SQL Server Information

If you do not want the management server to monitor a SQL Server instance, you can remove its information, as described in the following steps:

1. Select **Options > Protocol Settings > Storage Essentials > System Application Discovery Settings**.

To select a target, you must have at least one element designated as a server, workstation or desktop. If you see the message, "No Targets Currently Selected," change your element from unknown to either a server, workstation or desktop. Refer to the documentation for HP Systems Insight Manager.

2. Select a target, and then, click **Run Now**.
3. In the Database Information table, click the  button, corresponding to the SQL Server instance you do not want the management server to monitor.
4. Perform Discovery Data Collection to make the management server aware of your changes.

Monitoring Sybase Adaptive Server Enterprise

If you want to monitor Sybase Adaptive Server Enterprise you must:

- Create APPIQ_USER account on the database for Sybase
- Provide the database server name and port number
- Discover the application.

The required drivers for Sybase Adapter Server Enterprise were automatically installed along with the management server.

IMPORTANT: Before you begin these steps, make sure you purchased the module that lets you monitor Sybase Adaptive Server Enterprise. Contact your customer support if you are unsure if you purchased this module.

Step A - Create the APPIQ_USER account for Sybase

The management server accesses Sybase through the APPIQ_USER account. This account is created when you run the CreateSybaseAct.bat script on Microsoft Windows or CreateSybaseAct.sh on UNIX on the computer running the Sybase database you want to monitor. The account has create session and select dictionary privileges to be used with the management server.

Keep in mind the following:

- The script must run under SA user.
- Obtain the Sybase server name before you run the script
- Create APPIQ_USER account on Sybase Database you want to monitor.
- You should have already installed the database for the management server.
- Make sure you have all the necessary information before you begin the installation. Read through the following steps before you begin.

To create the APPIQ_USER account for the Sybase server:

1. Do one of the following:

- To run the script on IBM AIX, SGI IRIX, or Sun Solaris, log into an account that has administrative privileges, mount the CD-ROM (if not auto-mounted), and go to the /DBIQ/sybase/unix directory by typing the following:
cd /cdrom/cdrom0/DBIQ/sybase/unix
where /cdrom/cdrom0 is the name of the CD-ROM drive
- To run the script on Microsoft Windows, go to the \DBIQ\sybase\win directory on the CD-ROM.

IMPORTANT: You must complete the following steps.

2. Verify you have the password to the SA user account.
You are prompted for the password for this user account when you run the script.
3. Run the CreateSybaseAct.bat script on Microsoft Windows or CreateSybaseAct.sh script on the UNIX operating system on the computer with the Sybase database.
The script creates a user with login to master and select privilege on data dictionary tables on a managed Sybase instance.

NOTE: You can use a remote Sybase isql to run this script.

4. Type the Sybase instance name, which must be visible to the client, as the first input when running the script. The script prompts you for the name of the sybase server on which to create user for Sybase management packages and the password of the SA account.
5. Repeat the previous step for each Sybase server you want to manage.
This script does the following in order:
 - Creates the APPIQ_USER account.
 - Grant create session and select on dictionary tables privileges to APPIQ_USER enabling management server to view statistics for the Sybase server.

Removing the APPIQ_USER Account for Sybase

IMPORTANT: Before you remove the APPIQ_USER account for the Sybase databases on a host, make sure no processes are running APPIQ_USER for that Sybase database. The management server uses APPIQ_USER to obtain information about a Sybase database.

To remove the APPIQ_USER account for the Sybase databases on a host:

1. Do one of the following:
 - To run the script on IBM AIX, SGI IRIX, or Sun Solaris, log into an account that has administrative privileges, mount the CD-ROM (if not auto-mounted), and go to the /DBIQ/sybase/unix directory by typing the following:
cd /cdrom/cdrom0/DBIQ/sybase/unix
where /cdrom/cdrom0 is the name of the CD-ROM drive

- To run the script on Microsoft Windows, go to the \DBIQ\sybase\win directory on the CD-ROM.

IMPORTANT: You must complete the following steps.

2. Verify you have the password to the SA user account.
You are prompted for the password for this user account when you run the script.
3. Run the UninstallSybaseAct.bat script on Microsoft Windows or UninstallSybaseAct.sh script on the UNIX operating system on the computer with the Sybase database.
4. Type the name of the Sybase server.
5. Type the password for the SA account.
The account for APPIQ_USER is removed. The management server can no longer monitor the Sybase databases on this host.

Step B - Provide the Sybase Server Name and Port Number

You must provide the Sybase server name and port number for managing the Sybase database in the following steps:

To add information for discovering Sybase Adaptive Server Enterprise:

1. Select **Options > Protocol Settings > Storage Essentials > System Application Discovery Settings**.
To select a target, you must have at least one element designated as a server, workstation or desktop. If you see the message, "No Targets Currently Selected," change your element from unknown to either a server, workstation or desktop. Refer to the documentation for HP Systems Insight Manager.
2. Select a target, and then, click **Run Now**.
3. Click the **Create** button for the Database Information table.
4. In the **Host IP/DNS Name** field, type the IP address or DNS name of the host running Sybase.
5. You can leave the **Management IP/DNS Name** field blank. This field is for Oracle clusters. When you leave the **Management IP/DNS Name** field blank the management server automatically lists the DNS name or IP address of the host under the **Host IP/DNS Name** column and **Management IP/DNS Name** column.
6. In the **Server Name** field, type the Sybase database you want to monitor.
7. In the **Port Number** field, type the port that Sybase is using.
8. Select **SYBASE** from the Database Type menu.
9. Click **OK**.


IMPORTANT: Perform Discovery Data Collection for your inputs to take effect. See “[Step 3 - Discovering Applications](#)” on page 64.

Deleting Sybase Information

If you do not want the management server to monitor a Sybase instance, you can remove its information, as described in the following steps:

1. Select **Options > Protocol Settings > Storage Essentials > System Application Discovery Settings**.

To select a target, you must have at least one element designated as a server, workstation or desktop. If you see the message, “No Targets Currently Selected,” change your element from unknown to either a server, workstation or desktop. Refer to the documentation for HP Systems Insight Manager.

2. Select a target, and then, click **Run Now**.
3. In the Database Information table, click the  button, corresponding to the Sybase instance you do not want the management server to monitor.
4. Perform Discovery Data Collection to make the management server aware of your changes.

Monitoring Microsoft Exchange

To monitor Microsoft Exchange, you must make the management server aware of domain controller access. After information for controller access has been added, discover Microsoft Exchange and perform Discovery Data Collection. To save time, delay these steps until you have added the configurations for your other applications and hosts.

To monitor Microsoft Exchange, you must:

- Add information for Microsoft Exchange Domain Controller Access
- Discover the application.

Adding Microsoft Exchange Domain Controller Access

To obtain information about your Microsoft Exchange servers, you must provide the user name and password for at least a primary domain controller, in addition to a DNS name, as described in the following steps.

IMPORTANT: The hosts should recognize the management server by name, as a reverse look-up is required by operating system security as well as Microsoft Exchange.

To provide information about the Microsoft Exchange servers:

1. Select **Options > Protocol Settings > Storage Essentials > Global Application Discovery Settings**.


The information you provide for the primary domain controller and backup domain controller apply to all Microsoft Exchange servers you discover.

2. In the Microsoft Exchange Configuration section, click the **Edit** button.
3. Under the Primary Domain Controller section, perform the following steps:
 - a. In the **Host Name** field, type the fully qualified DNS name for the domain controller.
 - b. In the **User Name** field, type the user name for accessing the Microsoft Exchange server.
 - c. In the **Domain Password** field, type the corresponding password for accessing the Microsoft Exchange server.
 - d. In the **Verify Password** field, re-type the password for verification.
4. Under the Backup Domain Controller section, perform the following steps:
 - a. In the **Host Name** field, type the fully qualified DNS name for the domain controller.
 - b. In the **User Name** field, type the user name for accessing the Microsoft Exchange server.
 - c. In the **Domain Password** field, type the corresponding password for accessing the domain controller.
 - d. In the **Verify Password** field, re-type the password for verification.
5. Click the **OK** button.

IMPORTANT: You must discover the host running Microsoft Exchange. See "[Step 3 - Discovering Applications](#)" on page 64.

Deleting a Microsoft Exchange Domain Controller

To delete a Microsoft Exchange domain controller:

1. Select **Options > Protocol Settings > Storage Essentials > Global Application Discovery Settings**.
2. Click the  button, corresponding to the domain controller you want to remove.
3. Perform Discovery Data Collection (**Discovery > Details**) for your changes to take effect.

Step 3 - Discovering Applications

This step assumes you have already discovered your hosts and provided discovery information for your applications. To discover an application, do the following;

- Detect the application ("[Step A - Detect the Applications](#)" on page 64)
- Discovery Data Collection ("[Step B - Obtain Details](#)" on page 64)

Keep in mind the following:

- This section assumes you have already set up the discovery configurations for your applications as described in "[Step 2 - Setting Up Discovery for Applications](#)" on page 51.
- If DNS records for your Microsoft Exchange servers are outdated or missing, the discovery of Microsoft Exchange may fail because Microsoft Exchange is dependant on Active Directory, which is dependant on DNS. Since Active Directory is dependant on DNS, Active Directory replication and Active Directory lookups may fail or contain errors if DNS records are not accurate.

- The management server is unable to discover Oracle on a Windows host if the host is on a private network behind a Windows proxy. The management server can discover the Windows host through the Windows proxy, but the management server is not able to detect Oracle.

Step A - Detect the Applications

Use HP Systems Insight Manager to discover the applications. Refer to the HP Systems Insight Manager documentation.

Step B - Obtain Details

Obtain detailed information from the discovered applications as described in this section.

Keep in mind the following:

- Discovery Data Collection takes some time. You might want to perform this process when the network and the managed elements are not busy.
- If you want to enable File SRM for a host, make sure the File SRM option is selected.
- When you do Discovery Data Collection that includes an AIX host, the system log displays three SCSI errors (2 FSCSI error and 1 FCS error) per IBM adapter port. You can ignore these errors.
- You can quarantine elements to exclude them from Discovery Data Collection. See ["Placing an Element in Quarantine"](#) on page 46 for more information. Let us assume you want to discover all the elements in a discovery group, except for one. Perhaps the element you want to quarantine is being taken off the network for maintenance. You can use the quarantine feature to exclude one or more elements from discovery.
- If the management server unable to obtain information from an element during Discovery Data Collection as a result of a CIM Extension hanging, the management server places the access point where the CIM Extension is located in quarantine. The management server then moves onto getting details for the next element in the Discovery Data Collection table. These elements appear as missing until they are removed from quarantine. See ["Removing an Element from Quarantine"](#) on page 47 for information on how to remove an element from quarantine.

To obtain details:

1. Select **Options > Storage Essentials > Discovery > Run Discovery Data Collection**.
2. Verify the **File SRM** option is selected. The File SRM option appears for hosts that have the CIM Extension and an operating system that supports File SRM.

NOTE: If you plan to have File SRM scan a host, make sure you have 220 to 230 MB for each set of 1 million files.

3. Select the discovery group from which you want to Discovery Data Collection. If you are obtaining Discovery Data Collection for hosts for the first time, make sure **All Discovery Groups** is selected.

You can use discovery groups to break up getting Discovery Data Collection. For example, instead of Discovery Data Collection for all of the elements, you could specify that the management server gets the element details for only the elements in Discovery Group 1, thus,

saving you time. You add an element to a discovery group by modifying the properties used to discover the element. See ["Modifying the Properties of a Discovered Address"](#) on page 37.

4. Click the **Get Details** button.

IMPORTANT: If the management server cannot communicate with an application, it labels the application as "Discovered". The management server could find the application, but it could not obtain additional information about it.


5. Refer to the topic, for information about automating the gathering of Discovery Data Collection. If you run into problems with discovery, see ["Troubleshooting"](#) on page 501.

Changing the Oracle TNS Listener Port

The software uses port 1521 by default to communicate with the TNS Listener service on the Oracle server. If your port is different or you use multiple ports, you can assign a new port number.

IMPORTANT: The hosts should recognize the management server by name, as a reverse look-up is required by operating system security as well as the Oracle Transparent Name Substrate (TNS).

To change this port number or to add ports:

1. Select **Options > Protocol Settings > Storage Essentials > Global Application Discovery Settings**.
2. To assign a new port, click the **Create** for the **Oracle Information** table.
3. Type the new port number and click **OK**.
4. If necessary, click the  button to remove the old port number.
5. Verify all elements have been discovered by clicking the **Start Discovery** button.

See ["Troubleshooting Discovery and Discovery Data Collection"](#) on page 511 for more information.

Adding/Modifying Microsoft Exchange Domain Controller Access

To obtain information about your Microsoft Exchange servers, you must provide the user name and password for at least a primary domain controller, in addition to a DNS name, as described in the following steps.

IMPORTANT: The hosts should recognize the management server by name, as a reverse look-up is required by operating system security as well as Microsoft Exchange.

To provide information about the Microsoft Exchange servers:

1. Select **Options > Protocol Settings > Storage Essentials > System Application Discovery Settings**.

To select a target, you must have at least one element designated as a server, workstation or desktop. If you see the message, “No Targets Currently Selected,” change your element from unknown to either a server, workstation or desktop. Refer to the documentation for HP Systems Insight Manager.

2. Select a target, and then, click **Run Now**.
3. In the Microsoft Exchange Configuration section, click the **Edit** button.
4. Under the **Primary Domain Controller** section, perform the following steps:
 - a. In the **Host Name** field, type the fully qualified DNS name for the domain controller.
 - b. In the **User Name** field, type the user name for accessing the Microsoft Exchange server.
 - c. In the **Domain Password** field, type the corresponding password for accessing the Microsoft Exchange server.
 - d. In the **Verify Password** field, re-type the password for verification.
5. Under the **Backup Domain Controller** section, perform the following steps:
 - a. In the **Host Name** field, type the fully qualified DNS name for the domain controller.
 - b. In the **User Name** field, type the user name for accessing the Microsoft Exchange server.
 - c. In the **Domain Password** field, type the corresponding password for accessing the Microsoft Exchange server.
 - d. In the **Verify Password** field, retype the password for verification.
6. Click the **OK** button.
7. Verify all elements have been discovered by clicking the **Start Discovery** button.
8. Update the database with element changes. See [“Updating the Database with Element Changes”](#) on page 45.


See [“Troubleshooting Discovery and Discovery Data Collection”](#) on page 511 for more information.

Deleting a Microsoft Exchange Domain Controller

To delete a Microsoft Exchange domain controller:

1. Select **Options > Protocol Settings > Storage Essentials > System Application Discovery Settings**.

To select a target, you must have at least one element designated as a server, workstation or desktop. If you see the message, “No Targets Currently Selected,” change your element from unknown to either a server, workstation or desktop. Refer to the documentation for HP Systems Insight Manager.

2. Select a target, and then, click **Run Now**.
3. Click the  button, corresponding to the domain controller you want to remove.

Changing the Password for the Managed Database Account

The management server connects to database applications through the use of the APPIQ_USER account, an unprivileged account with read-only privileges. You can change the password the management server uses to connect to database applications, such as Oracle and Sybase. When

you change the password of APPIQ_USER, you must change the password of all database applications.

Keep in mind the following:

- Change the password in all database applications before you change the password through the user interface. The passwords must also match.
- You must enter a password in the **Password** and **Verify Password** fields.

To change the password:

1. Select **Options > Protocol Settings > Storage Essentials > System Application Discovery Settings**.

To select a target, you must have at least one element designated as a server, workstation or desktop. If you see the message, “No Targets Currently Selected,” change your element from unknown to either a server, workstation or desktop. Refer to the documentation for HP Systems Insight Manager.

2. Select a target, and then, click **Run Now**.
3. Click the **Change Password** button at the top of the page.
4. Verify you have already changed the password of the databases listed on this page.
5. Type a new password in the **Password** field.

The management server requires the password to have the following characteristics:

- a minimum of three characters
- starts with a letter
- contains only letters, numbers and underscores (_)
- does not start or end with an underscore (_)

6. Retype the password in the **Verify Password** field.
7. Click **OK**.

Obtaining Disk Drive Statistics from Engenio Storage Systems

IMPORTANT: Depending on your license, the ability to obtain disk drives statistics from Engenio storage systems may not be available. See the “List of Features” to determine if you have access to the additional statistics. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).

To obtain information about disk drive statistics from Engenio storage systems, you must install a CIM Extension on a host that can access the Engenio storage system. Ensure the proxy host has at least one LUN rendered by each controller of the array. Then, you must make the management server aware of that host, as described in the following steps:

1. Install the CIM Extension on a host that has access to the Engenio storage system.
2. Discover that host.
3. Select **Optimize > Storage Essentials > Performance Data Collection**.

4. Verify the **Data Collection** tab is displayed.
5. Click the **Start** button corresponding with the disk drive statistics for an Engenio storage system.
6. Set the date and time. See "[Setting the Date and Time for Performance Collectors](#)" on page 139.
7. Type a repeat interval and then select a unit of measurement from the drop-down menu.
The repeat interval determines how often the collectors gather the data.
8. Select a proxy host by clicking the **Browse** button.
9. Select a proxy host from the drop-down menu and then click **OK**.
The management server displays in the drop-down menu only hosts that are running a CIM Extension version 3.5 or later and have access to the corresponding Engenio storage system.

NOTE: You can always change the proxy host by returning to this page or by going to the **Properties** tab for an Engenio storage system. Double-click the Engenio storage system in System Manager. Click the **Properties** tab. Then, click the **Browse** button on the Properties tab.

10. Saving the proxy host may take time. When you are asked if you want to continue, click **OK**.
11. Click **OK** again to set the time for starting the collector.
12. If you do not see any hosts displayed verify you have the latest CIM Extension installed and running on a host that can access the Engenio storage system.

Assigning a File Extension in Netscape 7

Netscape 7 automatically assigns unknown files an HTML extension. To make Netscape 7 recognize the type of file, you must assign a file extension.

To assign a MIME type:

1. Click the download file link or button in the software.
2. Click the **Advanced** button in the lower-left corner.
3. In the **Description of type** field, delete the existing text and type a description of the file.
4. In the **File extension** field, delete the existing text and type the file extension.
5. Click **OK**.

The next time Netscape 7 sees the associated MIME type, it will assign the extension you typed in the **File Extension** field.

For example, in the following figure, the zip extension was assigned to a MIME type of application/unknown. The next time Netscape sees that MIME type, it will automatically assign the zip extension to the file.

6. Click **OK**.

4 Managing Security

IMPORTANT: Depending on your license, role-based security may not be available. See the “List of Features” to determine if you have access to role-based security. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).

This chapter describes the following:

- “[Managing User Accounts](#)” on page 77
- “[Managing Roles](#)” on page 83
- “[Managing Organizations](#)” on page 85

About the Security for the Management Server

The management server offers security based on roles and organizations. Role-based security determines access to certain functionality depending on the user account assigned to a role. Organizations determine if you can modify an element type, such as hosts. The management server ships with the Everything organization, which lets you modify all element types.

See the following topics for more information:

- “[About Roles](#)” on page 71
- “[About Organizations](#)” on page 74
- “[Planning Your Hierarchy](#)” on page 77
- “[Naming Organizations](#)” on page 77

About Roles

The management server ships with several predefined roles that are listed in the following table. These roles determine which components of the software a user can access.

For example, users assigned to the Help Desk role have access to Application Viewer and Event Monitoring for Storage Essentials, but not to System Manager, Provisioning Manager, Policy Manager, and Reporting. Likewise, users assigned to the domain administrator role have access to all of the features, as shown in the following table.

IMPORTANT: Roles only apply to features and elements in HP Storage Essentials. For example, assume you assigned a user to the Help Desk role in Storage Essentials. That user will have “view only” privileges only in Storage Essentials.

Table 7 Default Role Privileges

Feature	CIO (Chief Informa-tion Officer)	Domain Admini- strator	Storage Admini- strator	Server Admin- istrator	Applic- ation Admin- istrator	Help Desk
Application Viewer	X	X			X	X
System Manager	X	X	X	X	X	
Event Monitoring for Storage Essentials		X	X	X	X	X
Provisioning Manager		X	X			
Capacity Manager	X	X	X	X	X	
Policy Manager		X	X			
Chargeback Manager	X	X	X			
Business Tools	X	X	X			
Reporting	X	X	X	X	X	
Global Reporter	X	X	X			
File System Viewer		X		X		
Performance Manager	X	X	X	X	X	
Access CLI		X	X			
Custom Commands		X	X			
System Configuration		X				

Keep in mind the following:

- Users created in HP Systems Insight Manager are automatically placed in the SIMViewOnly role. This role does not allow users to access any of the features listed in [Table 7](#) on page 72. See “[Adding Users](#)” on page 78 for more information.

- Users with access to Global Reporter can view all the elements throughout the enterprise, including those on the server running Global Reporter. Grant access to Global Reporter only to those, who should be allowed to view all elements. You may want to disable this functionality for some users.
- If the System Configuration option is selected for a role, all users assigned to that role will have administration capabilities, as shown in the following list. If you do not want users belonging to that role to have those capabilities, do not assign the System Configuration option:
 - The ability to set organizations and roles
 - Schedule discovery
 - Find the CIM log level
 - Save log files, e-mail log files
 - Save the database, backup the database, and schedule a database backup
 - Configure Event Monitoring for Storage Essentials, File System Viewer and Performance Manager
 - Configure reports and traps
 - Set up the management server to send e-mail

Roles also restrict access to element properties, element records, and Provisioning Manager, as shown in the following table.

Table 8 Default Role Privileges with Elements

Role	Application	Host	Switch	Storage System	Tape Library	Others
CIO	View	View	View	View	View	View
Domain Administrator	Full Control	Full Control	Full Control	Full Control	Full Control	Full Control
Storage Administrator	View	View	Full Control	Full Control	Full Control	Full Control
Server Administrator	View	Full Control	View	View	View	View
Application Administrator	Full Control	View	View	View	View	View
Help Desk	View	View	View	View	View	View
SIMViewOnly	View	View	View	View	View	View

By selecting one of the following options, users belonging to that role are restricted access:

- **Full Control** - Lets you view and modify the record for the element (Asset Management tab) and perform provisioning if applicable.
- **Element Control** - Lets you view and modify the record for the element (Asset Management tab). Provisioning cannot be performed.
- **View** - Lets you only view element properties.

For example, if a user belongs to a role that only lets you view the element properties on storage systems, that user would not be allowed to perform provisioning on storage systems because their role does not have the **Full Control** option selected for storage systems. That same role could also have the **Full Control** option selected for switches, allowing the user to perform provisioning for switches. Thus, the user would not be able to provision storage systems, but the user would be able to provision switches.

You can modify roles and/or create new ones. For example, you can modify the Help Desk role so that the users assigned to this role can also view Reporting and modify servers.

About Organizations

IMPORTANT: Organizations only apply to elements in HP Storage Essentials. For example, assume you assigned a user to an organization containing only hosts. That user will be able to view only hosts in Storage Essentials; however, that user may be able to view all other elements in HP Systems Insight Manager.

You can specify which elements users can access. For example, you can specify that some users have only access to certain switches and hosts. However, these users must already be assigned to roles that allow them to see switches and hosts.

Users only assigned to the organization can see just the elements that belong to the organization. If users are assigned to more than one organization, they see all elements that belong to the organizations to which they are assigned. For example, assume you created two organizations: One called OnlyHosts that allowed access to only hosts and another called OnlySwitches that allowed access to only switches. If you assigned a user to OnlyHosts and OnlySwitches, they would have access to hosts and switches because those elements are listed in at least one of the organizations.

Organizations can also contain other organizations. An organization contained within another is called a child. The organization containing a child organization is called a parent. In the following figure, the BostonWebHosts organization contains two child organizations,

BostonWebHost_Windows and BostonWebHost_Solaris. BostonWebHosts is a parent because it contains two organizations.




Figure 5 Parent-Child Hierarchy for Organizations

If a child contains organizations, it is also a parent. Let's assume you add two organizations called BostonWebMarketing and BostonWebProduction to BostonWebHost_Windows.

BostonWebHost_Windows would become a parent because it now contains two organizations. It would also be a child because it is contained in BostonWebHosts.

Parent organizations allow access to all elements listed in their child organizations. For example, users assigned to the organization BostonWebHosts can access not only the elements in BostonWebHost_Windows, but also those in BostonWebHost_Solaris. This is because BostonWebHosts is a parent of the two child organizations.

The parent-child hierarchy for organizations saves you time when you add new elements. You need to add a new element only once. The change ripples through the hierarchy. For example, assume you add an element to BostonWebHost_Windows. Users not only assigned to BostonWebHost_Windows would see this addition, but also users assigned to any of the parent organizations containing BostonWebHost_Windows. For example, users assigned to BostonWebHosts would also see the addition because it contains BostonWebHost_Windows. Users, however, assigned to only BostonWebHost_Solaris would not see the addition.

A child organization can be in multiple parent organizations. As shown in the following figure BostonWebHosts and NYWebHosts are not only children of the WebHosts organization, but they are also children of the US East Coast organization. Assume you have a user that oversees all Web hosts in the company, you could assign them to the WebHosts organization. Users managing hosts and storage systems on the east coast would be assigned to the US East Coast organization, which is a parent of BostonWebHosts, NYWebHosts, and StorageSystems organizations. Assume an element is added to NYWebHost_Solaris, users assigned to one or more of the following organizations would see the addition:

- NYWebHost_Solaris
- NYWebHosts
- WebHosts
- US East Coast




Figure 6 Children in Multiple Organizations

When you remove an element from an organization, users belonging to that organization or to one of its parents can no longer access that element if it is not a member of any other organization. For example, assume an element named *MyHost* was not only a member of *BostonWebHost_Solaris*, but also had mistakenly become a member of *BostonWebHost_Windows*. If you remove *MyHost* from *BostonWebHost_Solaris*, users belonging to *BostonWebHost_Solaris* can no longer access the element. Users belonging to the following organizations would still see the element because the element is still a member of *BostonWebHost_Windows*.

- *BostonWebHosts*
- *WebHosts*
- *US East Coast*

Keep in mind the following:

- You cannot edit the Everything organization.
- If you don't want users to have access to all elements, remove them from the Everything organization.
- Users can access information about all elements from security pages. For example, if a user is assigned to an organization that has access to only hosts, the user can still see information relating to storage systems on the security pages.
- Because generic elements cannot be placed into organizations, they do not appear unless the user belongs to the Everything organization. You can tell if an element is generic by double-clicking the element in System Manager and then clicking the Properties tab. Look for its description on this page. For example, if it is listed as "Generic Host" the element is a generic host.
- Discovery lists (**Discovery** tab) are not filtered. Users can see all elements in the discovery lists regardless of their affiliation with an organization.

- Reports only display elements assigned to the user's organization, including child organizations. For example, if you attempt to view a Host Summary report and you do not have permission to access hosts through your organization, you are not given information about the hosts in the report. This is also true for e-mailing reports. Let's assume again you do not have permission to access hosts. The reports you e-mail will not contain information about hosts, including the host specific reports. If the users receiving your reports want to be able to view information about hosts, one of the following must happen:
 - The hosts in question must be added to your organization.
 - Someone else, who has the hosts in question already in their organization, must send the reports.

IMPORTANT: When adding a child to an organization, do not add the organization's parent as a child. For example, assume you created an organization named Child1 that has a parent organization named Parent1. When you are adding child organizations to Child1, do not select Parent1, as this creates a loop.

Planning Your Hierarchy

Before you begin creating organizations, plan your hierarchy. Do you want the hierarchy to be based on location, departments, hardware, software or tasks? Perhaps you want a combination of these options.

To help you with your task, create a table of users who manage elements on the network and the elements they must access to do their job. You might start seeing groups of users who oversee the same or similar elements. This table may help you in assigning users to the appropriate organizations.

Once you are done with planning your hierarchy, draw the hierarchy in an graphics illustration program, so you can keep track of which organizations are parents and children.

Create the child organizations first, then their parents. See the topic, "[Adding an Organization](#)" on page 85 for more information.

Naming Organizations

When you create an organization, give it a name that reflects its members. For example, you might want to use one or more of the following as a guideline:

- Type of elements that are members of the organization, such as switches, Sun Solaris hosts
- Location of the elements, such as San Jose
- Task, such as backup machines

You may find that it is easy to forget which containers are parents and children. When you name an organization, you might want to include a portion of the name of the dominant parent organization. For example, assume you have two types of Web hosts in Boston: Microsoft Windows and Sun Solaris. You might name the two children organizations BostonWebHost_Windows and BostonWebHost_Solaris and their parent, BostonWebHosts.

Managing User Accounts

This section discusses the following topics:

- ["Adding Users"](#) on page 78
- ["Editing a User Account"](#) on page 79
- ["Deleting Users"](#) on page 80
- ["Modifying Your User Profile"](#) on page 80
- ["Modifying Your User Preferences"](#) on page 80
- ["Viewing the Properties of a Role"](#) on page 82
- ["Viewing the Properties of an Organization"](#) on page 82

Adding Users

To access the management server, users must enter a user name and password.

NOTE: The user name and password should be alpha-numeric. They cannot exceed more than 256 characters. The user name cannot begin with a number.

You must create your user account in HP Insight Manager SIM, as described in the following steps:

1. Select **Options > Security > Users and Authorizations**.
2. Click the User tab.
3. Click **New**.
4. Provide the following information:

- **Login name** - Provide a user name.
- **Domain** - Provide the domain name of the server running Storage Essentials.

You do not need to provide additional information. For more information about the other options mentioned on this page, access the documentation accompanying HP SIM.

5. Click **OK**.


The new user is created.

IMPORTANT: New users can view the toolbars for Storage Essentials and not have enough privileges to use its features. You must grant users privileges so they can use not only view the features in the toolbar, but use them as well.

6. To authorize a user to use the features in Storage Essentials:
 - a. Click **New**.
 - b. In the **New Authorizations** table, select the user.
 - c. Select **Manually assign toolbox and system/system group authorizations**.
 - d. In the **Selected Toolbox(es)** section, select **HP SE Tools**.

- e. In the **Select Systems** list box, select the systems you want the user to be able to manage. Select CMS (Central Management Server) if you want to access information about the server running HP Systems Insight Manager.
 - f. Click **OK**.
7. You must restart Storage Essentials.

When Storage Essentials restarts, it contacts HP Systems Insight Manager for information about accounts that have been added or removed.
8. Access Storage Essentials through one of the menu options, such as **Options > Storage Essentials > Email Settings**.
9. In the upper-right corner, select **Security > Users**.

Notice that the users you created in HP Systems Insight Manager are put in the SIMViewOnly Role. This role does not allow them to access any of the features in Storage Essentials.
10. To enable users access to features in Storage Essentials, assign the user to a different role by doing the following:
 - a. Click the **Edit** button () corresponding to the user account you want to modify.
 - b. To change the role assign the user account, select a new role from the **Role** drop-down menu.


If you don't find any roles that fit your needs, you can create a new one, as described in "[Adding Roles](#)" on page 83.
 - c. Click **OK**.

Editing a User Account

Keep in mind the following:

- The "admin" account acts differently than the other accounts. You cannot add or remove organizations from the "admin" account. You cannot remove the Everything organization from the "admin" account. New organizations are automatically added to the "admin" account when they are created.
- This change takes effect immediately, even if the user is logged into the management server.

To modify a user account:

1. Access Storage Essentials through one of the menu options, such as **Options > Storage Essentials > Email Settings**.
2. In the upper-right corner, select **Security > Users**.
3. Click the **Edit** button () corresponding to the user account you want to modify.
4. To change the account name, type a new name for the user account in the **Name** field, for example: jsmith
This name becomes the user name for the account.
5. To change the name assigned to the user account, type a new full name for the account in the **Full Name** field.
This information is used to provide a correlation between an account name and a user.

6. To change the role assigned to the user account, select a new role from the **Role** drop-down menu.
7. To change the e-mail address listed, type a new e-mail address in the **E-mail** field.
8. To change the phone number listed, type the user's new phone number in the **Phone** field.
9. Change or remove information from the **Notes** field if necessary.
10. To change the organizations to which the user belongs, select or deselect the organizations from the table in the user interface.
The Everything organization is the default organization that lets users access all current and future elements.
11. Click **OK**.

Deleting Users

IMPORTANT: You cannot delete the admin account.

To delete a user account:

1. Access Storage Essentials through one of the menu options, such as **Options > Storage Essentials > Email Settings**.
2. In the upper-right corner, select **Security > Users**.
3. Click the corresponding **Delete** button (🗑️).
The user account is deleted.

Modifying Your User Profile

While you are logged into the management server, you can change the following aspects of your user profile:

- Full Name
- E-mail address
- Phone number

However, you are not allowed to modify the following information:

- Login Name
- Role
- Organization affiliation

If you want this information modified, contact your domain administrator. Your domain administrator makes these changes in the **Configuration** menu.

To modify your user profile, do the following:

1. Click the name of your account in the upper-left corner.



Figure 7 Changing Your User Profile

2. On the **User Profile** tab, modify one or more of the following:
 - **Full Name**
 - **E-mail address**
 - **Phone number**
3. When you are done with your modifications, click the **Save Changes** button.

Modifying Your User Preferences

Use the **User Preference** tab to modify your user preferences for System Manager, Element Topology, and Event Monitoring for Storage Essentials. The **User Preference** tab controls what is displayed for your user account.

To access the **User Preferences** tab:

1. Click the name of your account in the upper-left corner of Storage Essentials.



Figure 8 Accessing the User Preferences Tab

2. Click the **User Preferences** tab.

System Manager and Element Topology Preferences

To change the severity icons you view in System Manager and in the element topology, select a severity level from the **Display Severity icons with this severity level or higher drop-down** menu. See the topic, "[Viewing Event Status in the Topology](#)" on page 194 for more information about event status.

If you want events refreshed within a time period, select the **Refresh events automatically** field. Then, enter in minutes how often you want the event information on the screen updated. If this option is set to every five minutes, the management server refreshes the severity icons displayed in System Manager and the element topology every five minutes.

Event Monitoring for Storage Essentials Preferences

Use the following table as a guideline for changing your user preferences for Event Monitoring for Storage Essentials.

Table 9 Changing User Preferences for Event Monitoring for Storage Essentials

If you want...	Do the following...
To be reminded whenever you change a filter	Select the option, Always remind me to apply filters when I change them.
Events refreshed automatically	Select the option, Refresh events automatically. Then, enter how often in minutes you want events refreshed.

Table 9 Changing User Preferences for Event Monitoring for Storage Essentials (continued)

If you want...	Do the following...
Change the number of events displayed on each page	Select the number of events to appear on a page from the Number of Events combo box.
Change the severities to be included	<p>Select a severity level you want displayed in Event Monitoring for Storage Essentials from the Severities to be Included drop-down menu. See the topic, "Selecting a Severity Level" on page 333.</p> <p>If you want to customize the filter for the severity levels, click the Custom button. See the topic, "Customizing the Severity Level Filter" on page 334 for more information.</p>
Change the element types to be included	<p>Select the element types you want to be included from the Element types to be included drop-down menu. Events from these elements are displayed in Event Monitoring for Storage Essentials. See the topic, "Filtering Events by Summary Text" on page 335.</p> <p>If you want to customize the filter for the element types, click the Custom button. See "Filtering Events by Summary Text" on page 335.</p>

Warnings for Slow Systems Operations

By default, the management server warns you when it encounters issues with handling large amounts of data from storage systems, such as long load times.

If you do not want to be warned, clear the **Warn about slow storage system operations** option on the **User Preferences** tab. See "[Modifying Your User Preferences](#)" on page 80 for information on how to access the **User Preferences** tab.

Viewing the Properties of a Role

You can quickly determine which components a user can access by viewing the properties of the user's role.

To view the properties of a role:

1. Access Storage Essentials through one of the menu options, such as **Options > Storage Essentials > Email Settings**.
2. In the upper-right corner, select **Security > Users**.
 1. Click **Security > Users**.
 2. In the **Role** column, click the name of the role.

This page displays the following information:

- **Role Name** - The name of the role. This name appears in the users table (**Security > Users**)
- **Role Description** - A description of the role.

- **Access Level** - Determines how much access the user has to a type of element, such as hosts, storage systems, switches, and applications. See "[About the Security for the Management Server](#)" on page 71 for more information.
- **Access to the Storage Authority Components** - Determines which components in the management server the user can access.

To learn how to edit a role, see the topic, "[Editing Roles](#)" on page 84.

Viewing the Properties of an Organization

You can quickly determine which elements a user can access by viewing the properties of the user's organization:

1. Access Storage Essentials through one of the menu options, such as **Options > Storage Essentials > Email Settings**.
2. In the upper-right corner, select **Security > Users**.
3. In the **Organization** column, click the name of a organization.
4. To determine which elements are in a child organization, click the link of the child organization.
5. To learn more about an element, click the element's link.

This page displays the following information:

- **Name** - The name of the organization. This name appears in the users table (**Security > Users**)
- **Description** - A description of the organization
- **Organization Members** - Determines which elements the user can access. See "[About the Security for the Management Server](#)" on page 71 for more information.

To learn how to edit an organization, see the topic, "[Editing Organizations](#)" on page 87.

Managing Roles

This section discusses the following topics:

- "[Adding Roles](#)" on page 83
- "[Editing Roles](#)" on page 84
- "[Deleting Roles](#)" on page 85

Adding Roles

The management server ships with several roles. You can add roles to accommodate your organization. For example, you might want to add a role for quality assurance. See the topic, "[About the Security for the Management Server](#)" on page 71 for more information about roles and organizations.

IMPORTANT: The **Role Name** and **Description** fields do not accept special characters, except spaces and the following characters: \$, -, ^, ., and _

To add a role:

1. Access Storage Essentials through one of the menu options, such as **Options > Storage Essentials > Email Settings**.
2. In the upper-right corner, select **Security > Roles**.
3. Click the **New Role** button.
4. In the **Role Name** field, type a name for the role. For example: Quality Assurance.
The name can contain spaces, but it cannot be longer than 256 characters.
5. In the **Description** field, type a description for the role. For example: Role for those in quality assurance.
You cannot type more than 1024 characters in the **Description** field.
6. Select an access level for each element type:
 - **Full Control** - Lets you view and modify the record for the element (Asset Management tab) and perform provisioning.
 - **Element Control** - Lets you view and modify the record for the element (Asset Management tab).
 - **View** - Lets you view element properties.For example, if a user belongs to a role that only lets you view the element properties on storage systems, that user would not be allowed to perform provisioning on storage systems because their role does not have the **Full Control** option selected for storage systems. That same role could also have the **Full Control** option selected for switches, allowing the user to perform provisioning for switches. Thus, the user would not be able to provision storage systems, but the user would be able to provision switches.
7. Select the features you want a user to be able to access. For example, if you want a user to have access to System Manager, select System Manager from the list.
See "[Management Server Components](#)" on page 5 for more information about these features.
8. Click **OK**.

Editing Roles


The software lets you modify the default roles and/or the roles you have created. See the topic, "[About the Security for the Management Server](#)" on page 71 for more information about roles and organizations.

Keep in mind the following:

- You cannot edit the domain admin role.
- After you click the **OK** button in the Edit Role window, any users assigned to the role you edited are logged out of the management server. Users see the changes when they log back into the management server.

- The **Role Name** and **Description** fields do not accept special characters, except spaces and the following characters: \$, -, ^, ., and _

To edit a role:


1. Access Storage Essentials through one of the menu options, such as **Options > Storage Essentials > Email Settings**.
2. In the upper-right corner, select **Security > Roles**.
3. Click the **Edit** button ()
4. To edit the name of the role, change the name in the **Role Name** field.
The name can contain spaces, but it cannot be longer than 256 characters.
5. To edit the description of the role, change the description in the **Description** field.
You cannot type more than 1024 characters in the **Description** field.
6. To change the access level, change the options selected in the table.
 - **Full Control** - Lets you view and modify the record for the element (Asset Management tab) and perform provisioning.
 - **Element Control** - Lets you view and modify the record for the element (Asset Management tab).
 - **View** - Lets you view element properties.

For example, if a user belongs to a role that only lets you view the element properties on storage systems, that user would not be allowed to perform provisioning on storage systems because their role does not have the **Full Control** option selected for storage systems. That same role could also have the **Full Control** option selected for switches, allowing the user to perform provisioning for switches. Thus, the user would not be able to provision storage systems, but the user would be able to provision switches.
7. Select the features you want a user to be able to access. For example, if you want a user to have access to System Manager, select System Manager from the list.
See "[Management Server Components](#)" on page 5 for more information about these features.
8. Click **OK**.

Deleting Roles

IMPORTANT: A role cannot be deleted if it contains a user.

To delete a role:

1. Access Storage Essentials through one of the menu options, such as **Options > Storage Essentials > Email Settings**.
2. In the upper-right corner, select **Security > Roles**.
3. Click **Roles** from the drop-down menu.
4. Click the corresponding **Delete** button ()
The role is deleted.

Managing Organizations

This section discusses the following topics:

- ["Adding an Organization"](#) on page 85
- ["Viewing Organizations"](#) on page 87
- ["Editing Organizations"](#) on page 87
- ["Deleting an Organization"](#) on page 88
- ["Removing Members from an Organization"](#) on page 89
- ["Filtering Organizations"](#) on page 90

Adding an Organization

You can create new organizations to restrict access to certain elements. For example, assume you do not want the help desk to have access to elements belonging to a certain group. You could create an organization that does not allow access to those elements. Once you assign users to that organization, they would only be able to access the elements you specified.

See the topic, ["About the Security for the Management Server"](#) on page 71 for more information about roles and organizations.

Keep in mind the following:

- You cannot add organizations to any user with the Domain Administrator role, which has access to all organizations by default.
- Create child organizations first, then their parents.
- When adding a child to an organization, do not add the organization's parent as a child. For example, assume you created an organization named Child1, which is contained in a parent organization named Parent1. When you are adding child organizations to Child1, do not select Parent1, as this creates a loop.
- All discovered elements are accessible in Business Tools, regardless of a user's restrictions. For example, assume your account belongs to an organization that has only hosts as members. If you run the business tool Switch Risk Analysis, the management server still provides information about whether the switches are a risk in your environment.

To add an organization:

1. Access Storage Essentials through one of the menu options, such as **Options > Storage Essentials > Email Settings**.
2. In the upper-right corner, select **Security > Roles**.
3. Click the **New Organizations** button.
4. In the **Name** field, type a name for the organization.
The organization name can contain spaces, but it cannot be longer than 256 characters and it cannot contain the carot (^) symbol.
5. In the **Description** field, type a description for the organization.
You cannot type more than 1024 characters in the **Description** field.
6. Click the **Add or Remove Members** button to determine which elements the user will see.

7. To add elements, expand the Elements node in the tree. Then, do at least one of the following:
 - **Select all elements of a certain type** - Just select the node for that element type. For example, you can select all the hosts by just clicking the Hosts node in the left pane. The elements in that category appear in the Organization Members pane.
 - **Select individual elements** - Expand the Elements node. Then, expand the node for the element type, for example the Applications node. Select the elements you want to add to the organization. The selected elements appear in the Organization Members pane.
8. To add organizations, do one of the following:
 - **Select all organizations** - Select the Organizations (top-level node). The selected organizations appear in the Organization Members pane.
 - **Select individual organizations** - Expand the Organizations node. Then, select the organization. The selected organization appears in the Organization Members pane.

The organizations in the Organization Members pane are listed as child organizations because they are now contained within the organization you are creating. See the topic, "[About the Security for the Management Server](#)" on page 71 for more information.
9. Once you are done adding elements, click **OK** in the Add or Remove Organization Members window.
10. Once you are done adding the organization, click **OK** in the Add Organization window.

Viewing Organizations

The Setup Organizations page lists the organizations and their descriptions, in addition to the number of elements, users and child organizations assigned to each organization.

The number of elements field provides the total number of elements assigned to an organization, not including those within the child organization. An organization containing only child organizations displays 0 under the No. of Elements column; however, users assigned to that organization would have access to the elements assigned to its child organizations.



New Organization						
Name	Description	No. of Elements	Child Organizations	No. of Users	Edit	Delete
Everything	The default organization that includes all elements	106		2		
Boston_Hosts		3		0		

Figure 9 Viewing Organizations

Access the Setup Organizations page by clicking **Security > Organizations** in Storage Essentials.

You can access information about child organizations by clicking their link under the Child Organization column.

Editing Organizations


See the topic, "[About the Security for the Management Server](#)" on page 71 for more information about roles and organizations.

When elements are removed from an organization, users belonging only to that organization are no longer able to access the removed elements.

Keep in mind the following:

- Depending on your license, role-based security may not be available. See the "List of Features" to determine if you have access to role-based security. The "List of Features" is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).
- The **Name** and **Description** fields in the Edit Organization window do not accept special characters, except spaces and the following characters: \$, -, ^, ., and _
- The organization name can contain spaces, but it cannot be longer than 256 characters.
- You cannot edit the Everything organization.
- When adding a child to an organization, do not add the organization's parent as a child. For example, assume you created an organization named Child1, which is contained in a parent organization named Parent1. When you are adding child organizations to Child1, do not select Parent1, as this creates a loop.

To edit an organization:

1. Access Storage Essentials through one of the menu options, such as **Options > Storage Essentials > Email Settings**.
2. In the upper-right corner, select **Security > Roles**.
3. Click the  button.
4. To change the name of the organization, type a new name in the **Name** field.
The organization name can contain spaces, but it cannot be longer than 256 characters and it cannot contain the carot (^) symbol.
5. To change the description of the organization, type a new description in the **Description** field.
You cannot type more than 1024 characters in the **Description** field.
6. Click the **Add or Remove Elements** button.
7. Add and remove elements as described in the topics, "[Adding an Organization](#)" on page 85 and "[Removing Members from an Organization](#)" on page 89.
8. Once you are done adding or removing elements, click **OK** in the Add Organization page.
9. In the Edit Organization page, click **OK**.

Deleting an Organization

When an organization is removed, users assigned only to that organization are no longer able to access the elements in the removed organization. For example, assume you belong to two organizations, onlyHosts and onlySwitchesandHosts. The organization onlyHosts contains only hosts, and the organization onlySwitchesandHosts contains only switches and hosts. If you delete


the onlySwitchesandHosts organization, you will still have access to hosts because you still belong to the onlyHosts organization.

Keep in mind the following:

- You cannot remove the Everything organization, which is the default organization.
- You can only remove organizations with no users.

Depending on your license, role-based security may not be available. See the “List of Features” to determine if you have access to role-based security. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).

To delete an organization:


1. Click **Security > Organizations**.
2. Click the  button corresponding to the organization you want to remove.
The software removes the organization.

Removing Members from an Organization

When you remove an element from an organization, users belonging to that organization or to one of its parents can no longer access that element if it is not a member of any other organization. For example, assume an element named MyHost was not only a member of BostonWebHost_Solaris, but also had mistakenly become a member of BostonWebHost_Windows. If you remove MyHost from BostonWebHost_Solaris, users belonging to BostonWebHost_Solaris can no longer access the element. Users belonging to the BostonWebHost_Windows organization or to its parent would still see the element.


IMPORTANT: Depending on your license, role-based security may not be available. See the “List of Features” to determine if you have access to role-based security. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).


Use one of the following methods to remove an element from an organization:

- In the Edit Organization window, click the  button corresponding to the element or child organization you want to remove from the organization.
- In the Add or Remove Organization Members window, select the element or child organization you want to remove in the right pane and then click the **Remove from Organization** button.
Use this method if you want to add and remove elements from an organization.

Accessing the Edit Organization Window


To access the Add Organization window

1. Access Storage Essentials through one of the menu options, such as **Options > Storage Essentials > Email Settings**.
2. In the upper-right corner, select **Security > Roles**.
3. Click the **Edit** () button corresponding with the organization you want to edit.

4. In the Edit Organization window, click the  button corresponding to the element or child organization you want to remove from the organization.

Accessing the Add Elements to Organization Window

To access the Add Elements to Organization window:

1. Access Storage Essentials through one of the menu options, such as **Options > Storage Essentials > Email Settings**.
2. In the upper-right corner, select **Security > Roles**.
3. Click the **Edit** () button corresponding with the organization you want to edit.
4. In the Edit Organization window, click the **Add or Remove Members** button.
5. Select the elements and/or organizations you want to remove in the right panel and then click the **Remove from Organization** button.
6. Click **OK**.


Filtering Organizations


The management server provides a filtering feature that lets you designate which organizations are active in your view. For example, assume you belong to an organization name Hosts and this organization contains two organizations: "Windows Hosts" and "Solaris Hosts." If you want to view elements only in "Windows Hosts" and not in "Solaris Hosts" organizations, you could use the filtering feature to activate only the "Windows Hosts" organization.

Keep in mind the following:

- If you do not want to view an element, deselect all child organizations containing the element. You must also deselect all parent organizations containing the child organization that has that element. For example, assume you do not want to view all Solaris hosts and all Solaris hosts are in the "Solaris Hosts" organization. The "Solaris Hosts" organization is contained in the Hosts organization. You must deselect the "Solaris Hosts" organization and the Hosts organization if you do not want to see Solaris hosts.
- Users belonging to the Everything organization see all organizations on the management server in the Filtering Organizations window.
- If you do not select any organizations for filtering, you do not see any elements in the topology.

To filter organizations:

1. Access Storage Essentials through one of the menu options, such as **Tools > Storage Essentials > System Manager**.
2. In Storage Essentials, click the  button at the top of the screen, or click the link listing the organizations you can view.

1. Click the  button at the top of the screen, or click the link listing the organizations you can view.




Figure 10 Clicking the Organization Link

2. Deselect the organizations containing the elements you do not want to obtain information about. Assume you want to view only the elements in the "Windows Hosts" organization, you would select only "Windows Hosts." Let's assume you have a parent organization named "Hosts" that contains "Solaris Hosts" and "Windows Hosts." You would need to deselect "Solaris Hosts" and "Hosts." "Hosts" would need to be deselected because it contains organizations other than "Windows Hosts."

Links are displays for the organizations if you belong to a role that has System Configuration capability. To learn more about the contents of an organization, click its link.

3. Click **OK**.

You can now only obtain information about elements in the active organizations. The active organizations are listed in the link next to the filter button, as shown in the following figure.



Figure 11 Active Organization

5 Managing Licenses

To learn more about licensing, see the following topics:

- ["About the License"](#) on page 93
- ["Importing a License File"](#) on page 96
- ["Viewing Cumulative Licenses"](#) on page 96
- ["Viewing a Specific License"](#) on page 97
- ["Deleting a License"](#) on page 97

About the License

The management server restricts the number elements it manages through its license. It is important you keep your license up to date with the requirements of your network. The management server has several different types of license restrictions, as shown in [Table 10](#) on page 93.

Table 10 License Requirements

Type of Restriction	Description	Unit of Measurement
For File System Viewer	<p>The management server determines licensing for File System Viewer through terabytes (TB). When you purchased File System Viewer, you were given a number of TB you were allowed by the management server to monitor.</p> <p>The management server detects the number of TB that are being monitored on file servers and verifies that number is at or below the purchased amount.</p>	Terabytes (TB)
MALs	<p>Managed application licenses (MALs) for applications. A MAL is the number of discovered instances of the following:</p> <ul style="list-style-type: none">• Microsoft Exchange• Microsoft SQL Server• Oracle• Sybase Adaptive Server Enterprise	Number of MALs

Table 10 License Requirements (continued)

Type of Restriction	Description	Unit of Measurement
MAPs	The management software restricts the number of hardware elements it manages through the use of managed access points (MAPs) for hardware. A MAP is the sum of all storage access ports of all hardware elements that the management server manages. See Table 11 on page 94 for more information.	Number of MAPs
Raw NetApp Capacity	The Raw NetApp Capacity is the total disk capacity (unformatted capacity) of all discovered NetApp filers.	Terabytes (TB)

IMPORTANT: The management server updates the Current Usage Summary table every 24 hours after the first update. Elements the management server has discovered before the update are not reflected in the Current Usage Summary table. The time for the update is determined when the management server was first started. For example, the first update of the Current Usage Summary table occurs six hours after the management server was first started. The following updates occur every 24 hours. If the management server was started for the first time at noon, the first update of the Current Usage Summary table would occur at 6 p.m. All following updates would always occur at 6 p.m.

MAPs are determined as described in the following table.

Table 11 Determining Managed Access Points

Element	Managed Access Point
Hosts	The managed access points (MAPs) are the number of Fibre Channel ports with a minimum of one MAP. If a host has no Fibre Channel ports, the software assumes one MAP. The software does count direct attached storage, provided they are supported by the management server.
Switches	All ports on a switch are counted as MAPs.
Storage systems	The sum of all front facing ports. Storage systems with FA ports the software does not support, such as mainframe attached FICON, are still counted as MAPs. However, the management server does not count MAPs from storage systems it does not support. See the release notes for information about supported storage systems.

Example 1:

Assume you have the following environment:

- Brocade (two switches of 12 ports each, one switch of 16 ports) - Total 40 ports
- McDATA (one switch of 64 ports) - Total 64 ports
- Windows 2000 and Solaris Hosts (10 hosts with two Fibre Channel connection each) - Total 20 ports
- EMC Subsystem (one subsystem with 16 Fibre Channel ports) - Total 16 ports

The software calculates 140 MAPs in this environment.

Example 2:

Assume you have the same configuration above, and you added several devices to your network that the management server does not support. There are still 140 MAPs in this environment since the management server does not count the ports from devices it does not support.

Example 3:

Assume you have the same configuration as the first example with two Windows 2000 hosts that are directly attached to storage systems with no Fibre Channel (FC) connections with a total of 0 FC ports, as shown in the following figure:




Figure 12 An Example of Direct Attached Storage

The software calculates four MAPs (see the figure), since we assume one MAP for each host even though it has no Fibre Channel ports. The storage systems are counted since they are supported by

the management server. If you include the MAPs from the first example (140 MAPs), it brings the total to 144 MAPs.

Importing a License File

If you cannot find the license file you want to import or if you are interested in expanding your license for managing additional MAPs and MALs, contact your software representative.

To import a license file:

1. Click **Deploy > License Manager > Storage Essentials > Manage Storage Essentials Keys** in HP Systems Insight Manager.
2. Click the **Import License File** button.
3. Click the **Browse** button.
You are shown the file system of the computer being used to access the management server.
4. Select the license file.
5. Click **OK**.

Viewing Cumulative Licenses

Use view cumulative license feature to view the complete number of MALs and MAPs the management server supports at the current time. The software adds up the number of MALs and MAPs from the licenses and takes into account the expiration date.

IMPORTANT: You cannot modify the license file since it is encrypted. If you want to increase the number of MALs and/or MAPs, contact technical support.

To view cumulative licenses:

1. Click **Deploy > License Manager > Storage Essentials > Manage Storage Essentials Keys** in HP Systems Insight Manager.
2. Click the **View Cumulative Licenses** button.

The properties for the cumulative licenses are displayed:

Notice that each feature has a property that is either set to true or false. If a value for a property is set to true, your license lets you access that feature. Likewise, if the value is set to false, the license does not let you access that feature.

You can determine how many MAPs and/or MALs all of your licenses support by looking at the properties of the license file. The management software adds the number of allowable MAPs and MALs from all the license files and places the totals in this cumulative summary. MAPs and MALs that have expired will not be included in the totals. The MALs are split into three properties, LICENSE_FSRM_SIZE_TB, LICENSE_MAL_DATABASE, LICENSE_MAL_EXCHANGE. See the following list for the properties used for tracking MAPs and MALs:

- LICENSE_FSRM_SIZE_TB - The amount of space in Terabytes you are allowed for File System Viewer.

- `LICENSE_MAL_DATABASE` - The number of database application instances the management server is allowed to monitor, such as Oracle and Sybase Adaptive Server Enterprise.
- `LICENSE_MAL_EXCHANGE` - The number of Microsoft Exchange instances the management server is allowed to monitor.
- `LICENSE_MAPS` - The number of MAPs the management software can manage.

You can determine what you are currently using by looking at the **Current Usage Summary** table at the bottom of the page.

Viewing a Specific License

IMPORTANT: Do not manually edit the license. If you want to increase the number of MALs and/or MAPs, contact technical support.

To view a license for a type of user:

1. Click **Deploy > License Manager > Storage Essentials > Manage Storage Essentials Keys** in HP Systems Insight Manager.

2. Click the  button corresponding to the license you want to view.

The license name and file name is listed, along with its properties.

You can determine how many MAPs and/or MALs this license supports by looking at the properties in the license file. However, that can be misleading if you have other licenses that also provide support for MAPs and MALs. It is suggested you look at the cumulative licenses to obtain a total of the MAPs and MALs that are supported. See the topic, "[Viewing Cumulative Licenses](#)" on page 96 for more information about viewing cumulative licenses.


The MALs are split into three properties, `LICENSE_FSRM_SIZE_TB`, `LICENSE_MAL_DATABASE`, `LICENSE_MAL_EXCHANGE`. See the following list for the properties used for tracking MAPs and MALs:

- `LICENSE_FSRM_SIZE_TB` - The amount of space in Terabytes you are allowed for File System Viewer.
- `LICENSE_MAL_DATABASE` - The number of database application instances the management server is allowed to monitor, such as Oracle and Sybase Adaptive Server Enterprise.
- `LICENSE_MAL_EXCHANGE` - The number of Microsoft Exchange instances the management server is allowed to monitor.
- `LICENSE_MAPS` - The number of MAPs the management software can manage.

Deleting a License

IMPORTANT: Before you delete a license, make sure you have made a copy of it. If you delete the wrong license, you may lose access to certain features and/or access to the product. The management server saves the license files in the %JBOS4_DIST%\licenses directory.

To delete a license:

1. Click **Deploy > License Manager > Storage Essentials > Manage Storage Essentials Keys** in HP Systems Insight Manager.
2. Click the  button corresponding to the license you want to delete.

6 Configuring the Management Server

This chapter contains the following topics:

- "Trap Generation" on page 99
- "Setting Up E-mail Notification" on page 99
- "Managing Getting Discovery Details" on page 100
- "Managing Product Health" on page 102
- "Managing Logging" on page 105
- "Managing the Display of Events" on page 110
- "Managing File System Viewer" on page 113
- "Managing Reports" on page 118
- "Managing Performance Collection" on page 136

Trap Generation

You can configure the software so that events received by the system generate SNMP traps, which the software can send to another event monitoring system, such as Micromuse™ Netcool® Solutions or HP OpenView. The software allows up to five SNMP trap destinations.

The software provides an SNMP MIB you can compile into your existing enterprise framework. This MIB contains trap definitions so your enterprise framework can understand the traps. The MIB can be found in the `AppIQ-Traps.mib` file located in the `[Install_Dir]\Tools` directory and in the `Tools` directory on the CIM Extension CD-ROM, where `[Install_Dir]` is the directory into which you installed the management server.

To configure trap forwarding:

3. Select the **Enable Trap Generation** option.
4. In the **SNMP Community String** field, type the SNMP community string, which is used for filtering.
Important: If you enter a value in the **SNMP Community String** field, the SNMP agent must know the SNMP community string entered in the field to receive the SNMP trap.
5. Click the **Save** button.
6. In the **New SNMP Destination** field, type the IP address of the server running an SNMP agent.
7. Click the **Add** button.

Setting Up E-mail Notification

The management server provides e-mail notification for reports and policies. For example, you can set up the management server to notify you by e-mail when the amount of free space on a host becomes too low.

You must assign an SMTP server from which the management server can send its e-mail notifications.

To configure e-mail notification:

1. Click **Options > Storage Essentials > Email Settings** in HP Systems Insight Manager.
2. Click **E-mail Server** at the top of the screen.
3. (Required) Select **Enabled** to enable e-mail notification.
4. (Required) In the **Name** field, type the DNS name or the IP address of the Simple Mail Transfer Protocol (SMTP) server you want to use to send the e-mail notification.
5. (Required) In the **Port** field, type the port of the SMTP server you want to use to send the e-mail notification.
6. In the **User Name** field, enter a user name for the SMTP server.
7. In the **Password** field, enter a password for the SMTP server.
8. In the **Verify Password** field, re-type the password you entered previously.
9. (Required) In the **Sender** field, type the e-mail address of the sender.
This address is displayed in the **From** field in the e-mail.
10. If you want the replies to go to an e-mail address other than the e-mail address specified in the **Sender** field, type the e-mail address you want to receive replies in the **Reply to** field.
11. Click the **Save** button.

IMPORTANT: You should try sending a test e-mail.

To send a test e-mail:

1. In the **To** field, type an e-mail address.
The software verifies the address entered has a correct form.
2. In the **Subject** field, type a subject to distinguish from notification of a real event, for example:
Testing E-mail Notification
3. In the **Message** field, type a message, for example:
I'm just testing e-mail notification.
4. Click the **Send Test Message** button.

Managing Getting Discovery Details

You can schedule the management server to obtain discovery details at a specified interval. To learn more about scheduling details, see the following topics:


- ["Adding a Discovery Schedule"](#) on page 100
- ["Disabling a Schedule"](#) on page 102
- ["Editing a Schedule"](#) on page 102
- ["Removing a Schedule"](#) on page 102

Adding a Discovery Schedule

Schedule getting details when the network is not busy.

IMPORTANT: All collectors are stopped during include infrastructure details. This means that during getting SAN details information, such as for Performance Manager, is not updated.

To schedule discovery details:

1. Click **Options > Storage Essentials > Discovery > Schedule Discovery Data Collection** in HP Systems Insight Manager.
2. Click the **New Schedule** button.
3. In the **Name** field, type a name for your discovery schedule.
4. In the **Description** field, type a description for your discovery schedule.
5. Select one or more of the following:
 - **Include infrastructure details** - During include infrastructure details, the management server gathers detailed information about the SAN infrastructure and the process can be network intensive.
 - **Force Device Manager Refresh** - If you want the device managers for HDS and EMC Symmetrix storage systems to obtain the latest information whenever getting discovery details occurs. The management server obtains most of its information for HDS and EMC Symmetrix storage systems from their device managers. If the device managers do not have the latest information, the management server also displays the outdated information. See the following topics for more information: ["Excluding EMC Symmetrix Storage Systems from Force Device Manager Refresh"](#) on page 41.
6. Select the **Enable** option.
7. Click the calendar icon, .
8. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Click the date on which you want the next time getting discovery details to run. Today's date is highlighted in pink. Click **Set**.

The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.
9. In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - **Second(s)**
 - **Minute(s)**
 - **Hour(s)**
 - **Day(s)**
 - **Week(s)**
10. Click **Next**.
11. Select the discovery groups you want included in the discovery.

IMPORTANT: Only the elements in the discovery groups you select are included in discovery.

12. Click **Finish** when done.

You can always disable a schedule. To learn more about how to remove the schedule, see the topic “[Disabling a Schedule](#)” on page 102.


Disabling a Schedule

To disable a schedule for getting SAN topology details:

1. Click **Options > Storage Essentials > Discovery > Schedule Discovery Data Collection** in HP Systems Insight Manager.
2. Clear the **Enable** option in the SAN Topology Schedule section.
The schedule is disabled.


Editing a Schedule

To edit a schedule:

1. Click **Options > Storage Essentials > Discovery > Schedule Discovery Data Collection** in HP Systems Insight Manager.
2. Click the  button corresponding to the discovery schedule you want to modify.
3. If necessary, change the following properties:
 - Name
 - Description
 - Type of discovery
 - Schedule of the discovery
4. Click **Next**.
5. If necessary, change the discovery group you want assigned to the schedule.
6. Click **Finish**.

Removing a Schedule

To remove a schedule:

1. Click **Options > Storage Essentials > Discovery > Get Details** in HP Systems Insight Manager.
2. Click the  button corresponding to the discovery schedule you want to remove.
The schedule is removed.

Managing Product Health

The Product Health menu option in Storage Essentials helps you to monitor and manage the management server. Product Health lets do the following:


IMPORTANT: To obtain information from Product Health, you must have already discovered the management server and included it in Discovery Data Collection. See ["Step 1 - Discovering Your Hosts"](#) on page 49 for information on how to discover a host.

- **Disk Space Monitoring** - This feature keeps track of the management server's use of disk space. See ["Enabling Disk Space Monitoring"](#) on page 103.
- **Database Alert Log** - The Database Alert Log scans the management server for critical errors at a specified interval and displays the information in its own chart. This setup frees up Event Monitoring for Storage Essentials for monitoring other elements. See ["Enabling the Scanning of Critical Events of the Management Server Database"](#) on page 109.
- **Log Files** - You can view and download product logs, as described in
- **Scheduled RMAN backups** - This feature lets you schedule RMAN backups. See ["Scheduling RMAN Hot Backups of the Database"](#) on page 148. If you see the buttons on this page disabled, the product is set to no archive mode. See ["Setting the Archive Mode"](#) on page 155 for more information about changing the archive mode.
- **Advanced** - This feature lets you modify advanced settings so you can configure the product to run optimally in your environment. See ["Modifying Java Memory Settings"](#) on page 105 and ["Customizing Properties"](#) on page 105 for more information.

Enabling Disk Space Monitoring

You can configure the management server to monitor itself just as it would any other element. This feature lets you monitor the amount of free space the management server has left. The management server uses disk space for many of its operations, such as when it collects performance data, gathers element properties, generates events, and creates a backup.

IMPORTANT: To obtain information from Product Health, you must have already discovered the management server and obtained element details from it. See ["Step 1 - Discovering Your Hosts"](#) on page 49 for information on how to discover a host.

1. Make sure you have already obtained Discovery Data Collection for the management server. Discover the management server in the same manner as you would discover a host.
2. Click **Options > Storage Essentials > Manage Product Health** in HP Systems Insight Manager.
3. Click the calendar icon, .
4. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Click the date on which you want to run the monitoring. Today's date is highlighted in pink. Click **Set**.

The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.

5. In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - **Second(s)**
 - **Minute(s)**
 - **Hour(s)**
 - **Day(s)**
 - **Week(s)**
6. Click **Save Changes**.
7. To view the results of the monitoring, click the **Result** tab periodically.

Viewing the Results of Disk Space Monitoring

To make sure your management server does not run out of space, you should check the results for disk space monitoring.

IMPORTANT: The Results tab appears empty if the management server has not been included in Discovery Data Collection (**Discovery > Details**). To obtain information from Product Health, you must have already discovered the management server and obtained element details from it. See ["Step 1 - Discovering Your Hosts"](#) on page 49 for information on how to discover a host.

To access the results for disk space monitoring:

1. Click **Options > Storage Essentials > Manage Product Health** in HP Systems Insight Manager.
2. Select **Disk Space** in the tree.
3. Click the **Results** tab in the Disk Space window.

The following information is displayed:

- Data/Time
- Disk Capacity
- Free Space
- Database Files
- Archive Files
- RMAN Files
- Temp Tablespace

Advanced Settings

This section describes the following:

- ["Modifying Java Memory Settings"](#) on page 105
- ["Customizing Properties"](#) on page 105

Modifying Java Memory Settings

Do not modify the Java memory settings on Advanced page (**Options > Storage Essentials > Manage Product Health > Advanced**) unless instructed to do so by technical support. Incorrectly changing these settings could adversely impact the performance of the software.

Customizing Properties

The management server lets you modify its properties. These properties control a variety of functionality, such as the ability to specify the number of time-outs for switches. You can always view the default setting of the properties by accessing the Default Properties page. To customize properties:

1. Click **Options > Storage Essentials > Manage Product Health**.
2. Click **Advanced** in the Disk Space tree.
3. Click **Show Default Properties** at the bottom of the page.
4. Copy the commands you want to modify. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
5. Return to the Advanced page (**Options > Storage Essentials > Manage Product Health**). Then, click **Advanced** in the Disk Space tree).
6. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
7. Make your changes in the **Custom Properties** field. Make sure the property is not commented out by removing the hash (#) symbol in front of the property.
8. When you are done, click **Save**.
9. Restart the service for the management server for your changes to take effect.

While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Managing Logging

This section provides information about the following:

- ["Accessing the Log Files"](#) on page 106
- ["Downloading Logs to a File"](#) on page 106
- ["Downloading the User Audit Log"](#) on page 107
- ["Downloading the Discovery Summary Log"](#) on page 107
- ["Displaying a Log File in a Command Prompt Window"](#) on page 107
- ["Changing the Provider Log Level"](#) on page 108
- ["Enabling the Scanning of Critical Events of the Management Server Database"](#) on page 109
- ["Viewing the Results of Logging"](#) on page 110

Accessing the Log Files

You can obtain information about the software's and CIMOM's transactions in the log files, which are in the [Install_Dir]\logs directory, where [Install_Dir] is the directory into which you installed the management server. CIMOM is a component in the CIM management infrastructure that handles the interaction between management applications and providers. The log files may contain information that is difficult to understand. The following listing provides the logs files that would be the most useful to you. It is not a complete list of all the logs. Some logs are for internal use only.

- `appiq.log` - Provides information about the transactions in the software including web messages, EJB information, and general exceptions.
- `cimom.log` - Provides information about threads with CIMOM, such as provider transactions.

If you want to view all log files, you can save them in a zip file as described in the topic, ["Downloading Logs to a File"](#) on page 106.

Downloading Logs to a File

If you run into problems with the management server, use the "download logs to a file" feature to track the problem. This feature saves all the log files in a zip file, which is then stamped with the date and time (24-hour clock).

NOTE: Some of the log files are generated only when you run certain features. For example the `reports.log` is only generated when you run reports.

To save all logs to a file:

1. Click **Options > Storage Essentials > Manage Product Health** in HP Systems Insight Manager.
2. Select **Log Files** in the Product Health tree.
3. Click the **Download Logs** button.
4. When you are asked if you want to open or save the file, do one of the following:
 - **Microsoft Internet Explorer** - Click the **Save** button.
 - **Netscape 7** - Select the **Save this file to disk** option.The Downloading window appears.
5. Type a name for the *.zip file and select the directory to which you want to save the file. Keep in mind the following:
 - Netscape 7 assumes the file is an HTML file. If you are running Netscape, make sure the file type is selected to **All Files** and the file extension is zip. You can make Netscape 7 recognize the zip file next time, by clicking the **Advanced** button.
 - Make sure the .zip file is saved to a location other than the local disk drives of the management server.
6. Click the **Save** button in the Save As window.

Downloading the User Audit Log

You can determine who has been accessing your machine by viewing the user audit log.

To access the user audit log:

1. Click **Options** > **Storage Essentials** > **Manage Product Health** in HP Systems Insight Manager.
2. Click **Download User Audit Log**.
3. Save the zip file.
4. Unzip the zip file.
5. Open the text file in a text editor. If you are using NotePad, you may need to select **Format** > **Word Wrap**.

You may notice that information is displayed as follows:

```
[2005-05-09 09:22:24] INFO [admin/1000/computername.companyname.com] login
succeeded
```

where

- [2005-05-09 09:22:24] - The time and date the action occurred.
- INFO - Level of warning
- [admin/1000/computername.companyname.com] - The user name and DNS name of the computer used to log into the management server. In this case, the user logged in as admin from computername.companyname.com.
- login succeeded - the action that occurred.

Downloading the Discovery Summary Log

You can view status information from Discovery Data Collection by viewing the discovery summary logs, as described in the following steps:

1. Click **Options** > **Storage Essentials** > **Manage Product Health** in HP Systems Insight Manager.
2. Click **Download Discovery Summary Log**.
3. Save the zip file.
4. Unzip the zip file.
5. Open the `GAEDSummary.log` file in a text editor. If you are using NotePad, you may need to select **Format** > **Word Wrap**.

Displaying a Log File in a Command Prompt Window

The software ships with `tail.exe`, which can display and update a log file in a command prompt window. This utility is extremely helpful if you do not want to enable the option that lets the management server service interact with the desktop.

To display a log file in a command prompt window:

1. Open a command prompt window.

2. Go to the [Install_Dir]\logs directory by entering the following at the command prompt, where [Install_Dir] is the directory into which you installed the management server.

```
c:\>cd [Install_Dir]\logs
```
3. Enter the following at the command prompt window:

```
c:\[Install_Dir]\logs>tail -f appiq.log
```

 where appiq.log is the log file you want displayed in the command prompt window.
 The tail.exe utility checks the file for updates and appends them to text in the command prompt window.

Tail.exe is a program distributed under the GNU General Public License. See <http://www.gnu.org> for more information.

Changing the Provider Log Level

The management server obtains information from its discovered elements through providers, which communicate with the hardware interface. These providers provide by default superficial logging to the [Install_Dir]\server\appiq\logs\appiq.log file, where [Install_Dir] is the directory into which you installed the management server. You can change the level of logging provided by selecting the new level of logging from the **Provider Log Setting** combo box on the Product Health page (**Options > Storage Essentials > Manage Product Health** in HP Systems Insight Manager), selecting **Log Files** and then clicking **Apply**. Only one logging level option can be selected.

Use the table below as a guideline for the different options. Several of the options mention providers. A provider is software that gathers information from an element, such as a switch.

Table 12 Logging Levels


Log Level Option	Description	Use When You
Default Logging	Provides superficial logging	Do not want additional logging.
Detailed tracing of Brocade provider	Provides detailed logging of the providers used to gather information from the Brocade switch.	Are having difficulty obtaining information from a Brocade switch.
Detailed tracing of CIM Object Manager	Provides detailed logging of the infrastructure that handles the interaction between management applications and providers. The CIM Object Manager supports services such as event notification, remote access, and query processing.	Are having difficulty obtaining information from the CIM Object Manager. You are having difficulty obtaining information from event notification, remote access, and query processing.
Detailed tracing of CLARiiON provider	Provides detailed logging of the providers used to gather information from CLARiiON storage systems.	Are having difficulty obtaining information from CLARiiON storage systems.

Table 12 Logging Levels (continued)

Log Level Option	Description	Use When You
Detailed tracing of HDS provider	Provides detailed logging of the providers used to gather information from HDS storage systems.	Are having difficulty obtaining information from HDS storage systems.
Detailed tracing of HOST/SERVER provider	Provides detailed logging of the providers used to gather information from hosts/servers.	Are having difficulty obtaining information from host/servers.
Detailed tracing of all providers	Provides detailed logging of the providers, which gather information from the elements.	Are having difficulty obtaining information from more than one type of element.
Detailed tracing of SYMMETRIX provider	Provides detailed logging of the providers used to gather information from EMC Symmetrix storage systems	Are having difficulty obtaining information from EMC Symmetrix storage systems.

Enabling the Scanning of Critical Events of the Management Server Database

You can configure the management server to scan for only critical events occurring with the database for the management server at a specified time interval. The management server displays the results of these scans under **Options > Storage Essentials > Manage Product Health > Log Files** in HP Systems Insight Manager. To enable this type of scanning, do the following:

1. Click **Options > Storage Essentials > Manage Product Health** in HP Systems Insight Manager.
2. Select **Log Files** in the Product Health tree.
3. Click the calendar icon, .
4. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Click the date on which you want to run the next backup of the database. Today's date is highlighted in pink. Click **Set**.

The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.

5. In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - **Second(s)**
 - **Minute(s)**
 - **Hour(s)**

- **Day(s)**
- **Week(s)**

NOTE: The minimal interval you can schedule is one day. If you select **Hour(s)**, **Minute(s)** or **Second(s)**, you must enter an interval that equals more than a day. For example, if you select **Hour(s)**, you must enter 24 or more. Just as if you select **Minute(s)**, you must enter 1440.

6. Click the **Enable** option.
7. Click the **Save Schedule** button.
8. To view the results of the scanning, click the **Result** tab periodically.

Viewing the Results of Logging

You can view when an error occurred and at what time it was discovered by accessing the Results tab for logging.

To access the Results tab:

1. Click **Options > Storage Essentials > Manage Product Health** in HP Systems Insight Manager.
2. Select Log Files in the tree.
3. Click the Results tab in the Log Files pane.

The following is displayed:

- Scan Date/Time
- Error Occurred Time
- Error Description

Managing the Display of Events

This section describes the following:

- ["Controlling the Display of Cleared and Deleted Events"](#) on page 110
- ["Modifying the Clearing and Deletion Frequency"](#) on page 111
- ["Configuring the Clearing of Events"](#) on page 112
- ["Configuring the Deletion of Events"](#) on page 113

Controlling the Display of Cleared and Deleted Events

You can control how the management server displays events by modifying one or more of the following:

- **The Clearing and Deletion Frequency** - The frequency table determines how often the user interface in Event Monitoring for Storage Essentials removes events and marks events as "cleared." Events always display as they occur in the user interface.

Important: Events are only removed or marked “cleared” when their automatic delay time is completed. See the topics, “[Configuring the Clearing of Events](#)” on page 112 and “[Configuring the Deletion of Events](#)” on page 113.

Assume you set the clearing and deletion frequency to every 15 minutes with the initial time at 11:00 a.m. So every 15 minutes the management server checks for events marked for deletion and for clearing, and it updated the user interface accordingly. Let's assume Informational events are set to be cleared every hour, and an Informational event occurs at 12:20 p.m. Exactly after an hour has passed, the management server marks this event as “cleared” but the user interface is not updated because the frequency update of the user interface is every 15 minutes. If you looked at the Event Monitoring for Storage Essentials at 1:35 p.m. the event would be marked cleared. See “[Modifying the Clearing and Deletion Frequency](#)” on page 111 for more information.

- **The Clearing of Events** - You can determine how often events are marked “cleared.” See “[Configuring the Clearing of Events](#)” on page 112.
- **The Deletion of Events** - You can determine how often events are deleted. Events are deleted every two weeks by default. See “[Configuring the Deletion of Events](#)” on page 113.


Modifying the Clearing and Deletion Frequency

You can modify how often the user interface in Event Monitoring for Storage Essentials removes events and marks events as “cleared.” Events are still displayed as they occur in the user interface.

IMPORTANT: Events are only removed or marked “cleared” when their automatic delay time is completed. See the topics, “[Configuring the Clearing of Events](#)” on page 112 and “[Configuring the Deletion of Events](#)” on page 113.

Assume you set the clearing and deletion frequency to every 15 minutes with the initial time at 11:00 a.m. So every 15 minutes the management server checks for events marked for deletion and for clearing, and it updated the user interface accordingly. Let's assume Informational events are set to be cleared every hour, and an Informational event occurs at 12:20 p.m. Exactly after an hour has passed, the management server marks this event as “cleared” but the user interface is not updated because the frequency update of the user interface is every 15 minutes. If you looked at the Event Monitoring for Storage Essentials at 1:35 p.m. the event would be marked cleared. See “[Modifying the Clearing and Deletion Frequency](#)” on page 111 for more information.

To modify the clearing and deletion frequency:

1. Select **Options > Events > Storage Essentials > Event Configuration**.
2. Click the calendar icon, .
3. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Click the date on which you want to run the clearing and deletion frequency. Today's date is highlighted in pink. Click **Set**.

The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.

4. In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - **Second(s)**
 - **Minute(s)**
 - **Hour(s)**
 - **Day(s)**
 - **Week(s)**

NOTE: The higher the interval, the more demand there is on the management server.

5. Select the **Enable** option.
6. Click the **Save Changes** button.

Configuring the Clearing of Events

Depending on the severity of an event, the management server may mark the event as “clear” after 60 minutes. Events designated as Major and Critical are never marked as “clear.” You can change the time delay in clearing an event, and you can specify that the management server never marks an event as “clear”.

To help you in filtering events, you may want to have unimportant events marked as cleared rather than automatically deleted. Depending on how you have configured the deletion of events, you can view the “cleared” events at a later time.

See the following table for the default settings for clearing events.

Table 13 Default Settings for Clearing Events

Severity Level	Default Time Delay to Clear the Event (Hours)
Unknown	1
Informational	1
Warning	1
Minor	1
Major	Never
Critical	Never

To change the default time delay to clear an event, do the following:

1. Select **Options > Events > Storage Essentials > Event Configuration**.

2. Do one of the following:
 - If you never want an event of the specified severity level marked as “cleared,” select the **Never** option next to the severity level.
 - If you want to change the delay time in clearing an event, select one of the following units of measurement from the combo box and type the number in the adjacent field:
 - **Minutes**
 - **Hours**
 - **Days**
 - **Weeks**
3. Click the **Save Changes** button.

Configuring the Deletion of Events

The management server automatically deletes events after two weeks by default. You can specify for each severity level different time periods for deleting events. For example, you could modify the management server to delete events with the Information severity level every two days. You could also specify the management server to never delete events with the Critical severity level.

To change the default time delay to delete an event, do the following:

1. Select **Options > Events > Storage Essentials > Event Configuration**.
2. Do one of the following:
 - If you never want an event of the specified severity level automatically deleted, select the **Never** option under the Automatic Delete Delay column.
 - If you want to change the delay time in deleting an event, select one of the following units of measurement from the combo box and type the number in the adjacent field:
 - **Minutes**
 - **Hours**
 - **Days**
 - **Weeks**

For example, if you want events that are a week old deleted, you would type 1 and then select **Weeks** in the combo box in the Automatic Delete Delay column.

3. Click the **Save Changes** button.

Managing File System Viewer

To learn more about managing File System Viewer, see the following topics:

- [“Managing Collectors for File System Viewer”](#) on page 113
- [“Setting Data Collection Properties”](#) on page 115
- [“Scheduling File System Viewer Scan Checker and Scan Status”](#) on page 116
- [“Performing a Test File System Viewer Scan”](#) on page 118

Managing Collectors for File System Viewer

IMPORTANT: All collectors are stopped during “Discovery Data Collection.” This means that during “Discovery Data Collection” data, such as for File System Viewer, is not updated.

The management server uses File SRM data collectors to obtain information from a file server. File System Viewer provides information about a file server, such as the number of its directories. While File SRM data collectors are obtaining information from a file server, a second File SRM data collection cannot begin. Sometimes the management server cannot detect when the original data collection completed. For example, the management server may have been down when the data collection on the file server completed, and so the management server did not received notification of the completion of the data collection. In that case, all data collection is blocked on that file server because the management server assumes data collection on the file server is still occurring.

To resolve this blockage, do one of the following:

- Enable the File Scan Checker. This tool enables file servers that have their data collection blocked to be scanned again. The File Scan Checker removes references to File SRM data collections that started more than one hour ago so the management server no longer assumes a data collection on the file server is still occurring. The advantage of enabling the File Scan Checker is that once it is enabled, no manual intervention is required to keep file server scanning working. See “[Scheduling File System Viewer Scan Checker and Scan Status](#)” on page 116.
- Stop the collector for that file server and re-start it.

To manage collectors for File System Viewer:

1. Click **Reports > Storage Essentials > Data Collection > File SRM Data** in HP Systems Insight Manager.
2. Verify the **Data Collection** tab is displayed.

Table 14 About the Collectors for File System Viewer


Column Heading	Description
Element	The name of the element from which this collector gathers information.
Next Scheduled Run	Displays the date and time when the collector is scheduled to run.
Interval	The interval in minutes the collector runs.
Running?	Gives the status of the collector. Collectors that are running display a check mark in this column.
Edit	To edit the schedule for running a collector, click the Edit  button. Then set the date and time. See “ Setting Data Collection Properties ” on page 115 for more information.

Table 14 About the Collectors for File System Viewer (continued)

Column Heading	Description
Action	<p>Displays one of the following buttons:</p> <ul style="list-style-type: none">• Stop - Stops gathering information for which the collector is responsible. To stop more than one collector at once, select more than one collector and then click the Stop Collectors button. When a data collector is stopped, the corresponding reports will not contain recent data. In some cases, the reports may be blank.• Start - Starts the collector. When you start a collector, it begins gathering information for the corresponding file server. For example, assume you installed File System Viewer on a host and you want to view reports about file management, such as partition consumption. You must start the collector for that host.

NOTE: The management server does not collect information from mapped network drives on computers running Microsoft Windows.


Setting Data Collection Properties


The management server obtains information from collectors that gather information from elements at specified times. If you want the management server to monitor file servers, you must start the collector for the specified file server.

For example, assume you installed File System Viewer on a host and you want to view reports about file management, such as partition consumption. You must start the collector for that host.

IMPORTANT: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.





To set scanning properties:

1. Click **Reports > Storage Essentials > Data Collection > File SRM Data** in HP Systems Insight Manager.
2. Verify you are looking at the **Data Collection** page.
3. Do one of the following:
 - Click the **Start** button corresponding to the file server you want to monitor.
 - Click the **Edit** button () corresponding to the collector you want to modify.
4. Under **Local Volumes**, select the local volumes you want to monitor.
5. Under **External Volumes**, select the external volumes you want to monitor.

6. Under **Network Volumes**, select the network volume you want to monitor. Network volumes are only available on Solaris file servers.
7. If you want more than one process to be used for gathering information, change the number in the **Thread** field. For example, if you type 2 in the Thread field, two processes are used to collect the information. The more processes you assign to collecting information, the faster the information is collected. If you assign a large number of processes, you may slow down the management server.
8. Click the calendar icon, .
9. Enter the time in the time field. Make sure the time resembles a 24-hour clock, for example 22:00 for 10 p.m.
10. Click the date.

The date is highlighted in pink.

You can navigate the calendar as follows:

-  - Displays the same month in the previous year
-  - Displays the previous month
-  - Displays the next month
-  - Displays the same month in the following year

11. When you are done, click the **Set** button.

The management server sets the time.

In the following figure, the calendar sets the date and time at 1:56 p.m. February 13, 2003.

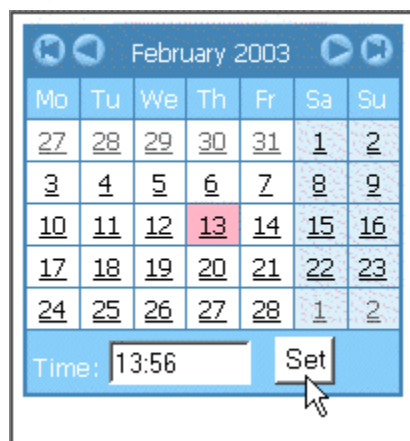


Figure 13 Selecting the Date and Time from the Calendar

12. Click **OK**.

Scheduling File System Viewer Scan Checker and Scan Status

The File Scan Checker tool enables file servers that have their data collection blocked to be scanned again. The File Scan Checker removes references to File SRM data collections that started

more than one hour ago so the management server no longer assumes a data collection on the file server is still occurring.

The advantage of enabling the File Scan Checker is that once it is enabled, no manual intervention is required to keep data collection on the file server working. See the steps in this section on how to enable and schedule the File Scan Checker.

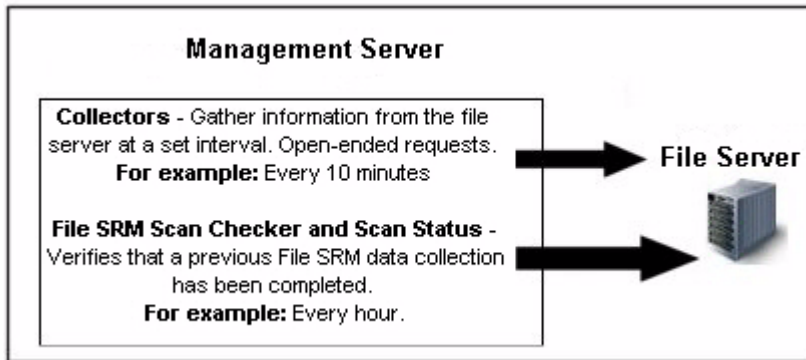


Figure 14 About the File System Viewer Scan Checker

To enable the File System Viewer scan checker and set its schedule:

1. Click **Reports > Storage Essentials > Data Collection > File SRM Data** in HP Systems Insight Manager.
2. Click the **Scan Status** tab.
3. Click the calendar icon,
4. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Today's date is highlighted in pink. Click the date you want the scan to run. Click **Set**.

The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.

5. In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - **Seconds(s)**
 - **Minute(s)**
 - **Hour(s)**
 - **Day(s)**
 - **Week(s)**
6. Select the **Enable** option.
7. Click the **Save Changes** button.

You can always disable the schedule by clearing the **Enable** option.

Performing a Test File System Viewer Scan

If you have changed your network configuration and you are not sure the management server can still scan the server running File System Viewer, you can perform a test File System Viewer scan. During this scan, the management server verifies that it can access the server and perform a high level File System Viewer scan.

IMPORTANT: This feature only works with CIM Extensions builds 4.0 and later.

To perform a test File System Viewer scan:

1. Click **Reports > Storage Essentials > Data Collection > File SRM Data** in HP Systems Insight Manager.
2. Click the **Test** button in the Test Scan column for the corresponding file server.
After you click the **Test** button, the management server attempts to access the server and perform a test File System Viewer scan.

Managing Reports

To learn more about managing reports, see one of the following topics:

- ["Architectural Overview of Report Views and Report Cache Refresh"](#) on page 118
- ["Managing Collectors for Reports"](#) on page 121
- ["Setting the Date and Time for Report Collectors"](#) on page 123
- ["Viewing Scheduled E-mail Deliveries for Reports"](#) on page 124
- ["Editing E-mail Schedules for Reports"](#) on page 125
- ["Viewing Data Aging Statistics for Reports"](#) on page 126
- ["Scheduling Report Cleanup"](#) on page 127
- ["Refreshing the Report Cache"](#) on page 128
- ["Setting Up Global Reporter"](#) on page 130

Architectural Overview of Report Views and Report Cache Refresh

Management server reports are based on report views (Materialized views) which are data snapshots of the management server schema at a certain time. The management server has a refresh process to refresh the report views. You can change the frequency of how often the report

The management server has two types of views for its reports. During a report cache refresh these views are updated. You can check the status of these views as described in this section:

- **MVIEWCORE_STATUS** - This table keeps track of the refresh status of the core views. The core views are the views starting with `mvc_` and `mvca`, as shown in [Table 58](#) on page 364. To obtain detailed information of **MVIEWCORE_STATUS**, see [Table 108](#) on page 397.
- **MVIEW_STATUS** - This table keeps track of the refresh status of the regular views, which are views starting with `mv_`, as shown in [Table 58](#) on page 364. To obtain detailed information of the **MVIEW_STATUS**, see [Table 109](#) on page 397.

To query the **MVIEW_STATUS** table:

1. Enter the following from the command prompt:

```
Sqlplus appiq_system/password
where password is the password for the appiq_system account.
```

2. Enter the following at the command prompt:

```
Sql>col lastRefresh format a30
This command ensures the data is displayed in a readable format.
```

3. Enter the following at the command prompt on one line with a space between the closing parenthesis and `lastRefresh`:

```
Sql> select mview_name, to_char(last_refresh_time,'mm/dd/yyyy hh24:mi:ss')
lastRefresh,
```

4. Enter the following at the command prompt:

```
status
```

5. Enter the following at the command prompt:

```
from mview_status
```

6. Enter the following at the command prompt:

```
order by 2;
```

To query the **MVIEWCORE_STATUS** table:

1. Enter the following from the command prompt:

```
Sqlplus appiq_system/password
where password is the password for the appiq_system account.
```

2. Enter the following at the command prompt:

```
Sql>col lastRefresh format a30
This command ensures the data is displayed in a readable format.
```

3. Enter the following at the command prompt on one line with a space between the closing parenthesis and `lastRefresh`:

```
Sql> select mview_name, to_char(last_refresh_time,'mm/dd/yyyy hh24:mi:ss')
lastRefresh,
```

4. Enter the following at the command prompt:

```
status
```

5. Enter the following at the command prompt:

```
from mviewcore_status
```

6. Enter the following at the command prompt:

```
order by 2;
```
7. Enter the following at the command prompt on one line with a space between the closing parenthesis and lastRefresh:

Managing Collectors for Reports

The management server uses statistics collectors to gather information for reports. To view a report you must have its corresponding collector running.

IMPORTANT: All collectors are stopped during “Discovery Data Collection.” This means that during “Discovery Data Collection” data, such as for Reports is not updated.

To view collectors for reports, click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager.

Table 15 About Collectors for Reports


Column Heading	Description
Element	The name of the element from which this collector gathers information.
Element Type	The type of element from which the collector gathers information.
Reports	The reports the collector is responsible for providing information.
Next Scheduled Run	Displays the date and time when the collector is scheduled to run.
Interval	The interval in minutes the collector runs.
Running?	Gives the status of the collector. Collectors that are running display a check mark in this column.
Edit	<p>To edit the schedule for running a collector, click the Edit () button. Then set the date and time. See “Setting the Date and Time for Report Collectors” on page 123 for more information.</p> <p>Important: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the first day of the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.</p>


Table 15 About Collectors for Reports (continued)

Column Heading	Description
Action	<p>Displays one of the following buttons:</p> <ul style="list-style-type: none">• Stop - Stops the collector. The corresponding reports display information only gathered previously. See "Stopping Collectors" on page 122.• Start - Starts the collector. When you start a collector, it begins gathering information for its corresponding reports. See "Starting Collectors" on page 122.

Starting Collectors

IMPORTANT: After you click **OK**, the collector starts immediately if the date and time you set has not passed. If the set time has passed, the collector starts two minutes after you click **OK**.

To start a collector:

1. Click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager.
2. Click the **Start** button corresponding to the collector you want to start.
3. Click the calendar icon, .
4. Set the date and time as described in the topic, "[Setting the Date and Time for Report Collectors](#)" on page 123.
5. When you are done with setting the date and time, enter an interval for the collector to gather information in minutes.

IMPORTANT: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.

6. Click **OK**.
7. To start more than one collector at once, select more than one collector and then click the **Start Collectors** button.

Stopping Collectors

When you stop a collector, the management server stops gathering information for which the collector is responsible. For example, if a collector is not running, its corresponding reports are no longer receiving information to display. One of the following occurs:

- If there was originally no information gathered for the report, no data appears in the report.

- If information was previously gathered for the report, old data appears in the report.

To stop a collector.



1. Click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager.
2. Click the **Stop** button corresponding to the collector you want to stop.
The collector stops gathering information for its corresponding reports.

To stop more than one collector at once, select more than one collector and then click the **Stop Collectors** button.

Setting the Date and Time for Report Collectors

IMPORTANT: After you click **OK**, the collector starts immediately if the date and time you set has not passed. If the set time has passed, the collector starts two minutes after you click **OK**.

To set the date and time, do the following:

1. Click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager.
2. Click the **Stop** button corresponding to the collector you want to stop.
3. Click the **Edit** button, .
4. Click the calendar icon, .
5. Enter the time in the time field. Make sure the time resembles a 24-hour clock, for example 22:00 for 10 p.m.
6. Click the date.

The date is highlighted in pink.

You can navigate the calendar as follows:



- Displays the same month in the previous year



- Displays the previous month



- Displays the next month



- Displays the same month in the following year

7. When you are done, click the **Set** button.

The management server displays the date and time.

IMPORTANT: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.

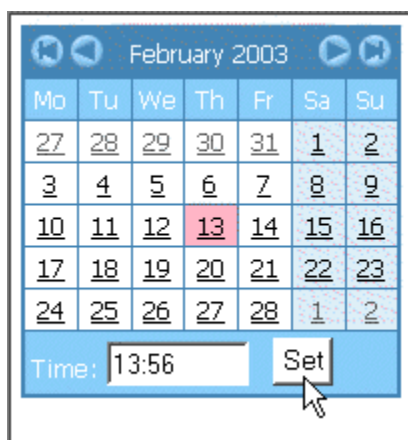


Figure 16 Selecting the Date and Time from the Calendar

In the previous figure, the calendar sets the date and time at 1:56 p.m. February 13, 2003.

Viewing Scheduled E-mail Deliveries for Reports

The Scheduled Deliveries tab displays all e-mail schedules for reports. An e-mail schedule instructs the management server to send a particular report to one or more recipients at a regular interval. For example, you could create an e-mail schedule that sends a host utilization report to your boss on a weekly basis. To learn more about creating e-mail schedules, see [“Adding an E-mail Schedule for a Report”](#) on page 344.



To view all e-mail schedules:

1. Click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager.
2. Click the **Scheduled Deliveries** tab at the top of the screen.
All e-mail schedules for reports are displayed.

Table 16 Viewing System Wide E-mail Schedules

Column Name	Description
Recipient	The person who receives the report.
Subject	The subject of the e-mail.


Table 16 Viewing System Wide E-mail Schedules (continued)

Column Name	Description
Format	The format of the report sent: <ul style="list-style-type: none">• PDF• Microsoft EXCEL• XML
Scheduled By	The user who scheduled the report.
Edit	Click the  button to edit a schedule. See "Editing an E-mail Schedule for a Report" on page 347 for information about the options displayed in this window.
Delete	Click the  button to remove the corresponding schedule. To delete multiple schedules, select the schedules you want to delete in the far left column. To quickly select all schedules, select the box to the left of the Recipient column. Then, click the Delete button above the table.

Editing E-mail Schedules for Reports

IMPORTANT: Schedule to send your reports soon after a report cache refresh. The reports display data that is in the report cache. If the report cache contains old data, the reports you send by e-mail will also show old data. The reports are refreshed every six hours by default. For example, assume you added an e-mail schedule that sends a report daily at 7 a.m. Also, assume you scheduled your report cache refreshes to take place daily at 8 a.m. Your reports will most likely show outdated data. It would make more sense to schedule your report cache refresh at 7 a.m. and then schedule to send your reports soon afterwards. See ["Scheduling a Report Cache Refresh"](#) on page 129.

To edit an e-mail schedule:

1. Click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager.
2. Click **Scheduled Deliveries** at the top of the screen.
3. Under the Edit column, click the **Edit** () button.
4. In the **To** field, change the recipient's e-mail address.
5. In the **Subject** field, change the subject of the e-mail.
6. In the **Message** field, change a message describing the report.

If you are e-mailing reports in bulk, you might want to let users know the e-mail is being sent by an automated process. You might also want to provide an e-mail address for users to provide feedback, for example:

This e-mail and its attached report are generated automatically. If you would like to change how often the report is sent to you or you want to be taken off the list, please contact `username@companyname.com`.

7. From the **Format** drop-down menu, select one of the following formats:
 - **PDF** - Requires the use of Adobe Acrobat, which can be downloaded for free from <http://www.adobe.com>.
 - **Excel** - Requires the use of Microsoft Excel.
 - **XML** - Requires the user has an understanding of XML.
8. In the **Time to Run** field, type the time you want to send the report. This time must be entered in the 24-hour format. For example, if you want a report sent at 2:15 p.m., you would type 14:15 in **Time to Run** field.
9. Select one of the following options to determine how frequently you want to send the report.
 - **Daily** - If you selected daily, select how frequently you want the management server to send the report.
 - **Everyday** - The report is sent everyday.
 - **Weekday** - The report is sent only Monday through Friday.
 - **Everyday for a specified number of days** - Fill in the number of days you want the report to be sent daily. After the specified number of days, the report is no longer sent. For example, you could use this feature to send reports to someone's replacement while the person is away on vacation.
 - **Weekly** - If you selected weekly, use the **Frequency** drop-down menu to select the day of the week on which you want the report sent.
 - **Monthly** - If you selected monthly, select the time during the month you want the report sent.
10. Click the **OK** button.

Viewing Data Aging Statistics for Reports

Data Aging includes Data Rollup and Garbage Collection. Data Rollup controls how often a set of data is summarized. For example, hourly data is rolled into the daily table periodically. Garbage Collection refers to how long a set of data is preserved before it is permanently removed from the database.

The settings on this page control data aging for both reports and performance statistics.


IMPORTANT: Do not modify the data on this page unless instructed by customer support. Changing them incorrectly can adversely affect the management server.

To view data aging statistics:

1. Click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager.
2. Click the **Data Aging** tab at the top of the screen.

Data aging statistics are displayed in the table.

IMPORTANT: Perform the following steps **ONLY** if customer support has instructed you to modify one of the collectors on the page:

3. Click the **Edit** () button.
4. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Today's date is highlighted in pink. Click **Set**.
The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd format.
If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.
5. Click the **Enable** option.

NOTE: If you are not allowed to disable the collector, the Enable option is unavailable.


6. To change the repeat interval, In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - **Second(s)**
 - **Minute(s)**
 - **Hour(s)**
 - **Day(s)**
 - **Week(s)**
7. (Garbage Collection only) To change how long the data is preserved, type an interval in the **Preserve** field, and then select one of the following from the drop-down menu to the right of the field:
 - **Second(s)**
 - **Minutes**
 - **Hours**
 - **Days**
 - **Weeks**
8. Click the **OK** button.

Scheduling Report Cleanup

Temporary files are created when you use Reporting to create reports. The report cleanup does not delete any data. It is used primarily for deleting temp files related to report management, such as when a report is viewed in the PDF or Microsoft Excel format.

You can schedule how often these files are removed, as described in the following steps:

1. Click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager.

2. Click the **Report Cleanup** tab at the top of the screen.
3. Click the calendar icon, .
4. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Click the date on which you want to run the next report cleanup. Today's date is highlighted in pink. Click **Set**.

The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.
5. In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - **Second(s)**
 - **Minute(s)**
 - **Hour(s)**
 - **Day(s)**
 - **Week(s)**
6. Click the **Enable** option.
7. Click the **Save Changes** button.

Refreshing the Report Cache

The management server gathers information for reports from the database every six hours by default. The management server stores this information in its report cache and displays it when a report is requested. If you are seeing outdated information in the report, you can refresh the report using one of the following methods:

- **Refresh the report cache now.** - See "[Refreshing the Report Cache Immediately](#)" on page 128.
- **Schedule the report cache to be refreshed.** - See "[Scheduling a Report Cache Refresh](#)" on page 129.

Refreshing the Report Cache Immediately

The **Refresh Now** button lets you update the report cache immediately.

Keep in mind the following:

- If Discovery Data Collection is occurring, wait for it to finish before clicking the **Refresh Now** button. This technique ensures the database is completely updated and thus, your reports will be as accurate as possible. Discovery Data Collection collects the latest data. When you refresh the report cache, the management server transfers the information collected from Discovery Data Collection and transfers it to the report cache.
- While the report cache is being refreshed, reports display no data.

To refresh the report cache now:

1. Click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager.
2. Click the **Report Cache** tab at the top of the screen.
3. Click **Refresh Now**.

Scheduling a Report Cache Refresh


The management server lets you schedule the refreshing of the report cache.

Keep in mind the following:

- The higher the frequency of the report cache interval, the more stress you put on the management server. A very frequent report cache interval, such as every 10 minutes, could hurt the response time of the management server to perform other tasks.
- Make sure you schedule refreshing the report cache after Discovery Data Collection. This technique ensures the database is completely updated and thus, your reports will be as accurate as possible. Discovery Data Collection collects the latest data. When you refresh the report cache, the management server transfers the information collected from Discovery Data Collection and transfers it to the report.
- While the report cache is being refreshed, reports display no data.

If you find you are still viewing old information regarding elements on the network, you may need to perform Discovery Data Collection. It is best to perform Discovery Data Collection at regular intervals. See "[Adding a Discovery Schedule](#)" on page 100.

To schedule a report cache refresh:

1. Click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager.
2. Click the **Report Cache** tab at the top of the screen.
3. Click the calendar icon, .
4. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Click the date on which you want to run the next report cleanup. Today's date is highlighted in pink. Click **Set**.

The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.

5. In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - **Second(s)**
 - **Minute(s)**
 - **Hour(s)**
 - **Day(s)**
 - **Week(s)**

6. Click the **Enable** option.
7. Click the **Save Changes** button.

Setting Up Global Reporter

IMPORTANT: Depending on your license, Global Reporter may not be available. See the “List of Features” to determine if you have access to Global Reporter. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).

Global Reporter lets you gather data for global reports from multiple management servers. For example, let's assume you have three management servers: one in London, one in Tokyo, and one in New York City. You are located in New York City and you want to obtain data in your global reports from the management servers in London and Tokyo, as well as from your own management server in New York City. Through the use of Global Reporter, you can gather data from all three sites.

When you set up Global Reporter, the management server pulls the data from the local database views at these sites through the database link. The flow of the data is shown by the arrows in the following figure. A global reporting view on the server running Global Reporter contains

information in the database that can be used for local and global reports. The management server in New York is referred to as a Global Reporter server because it has Global Reporter enabled.

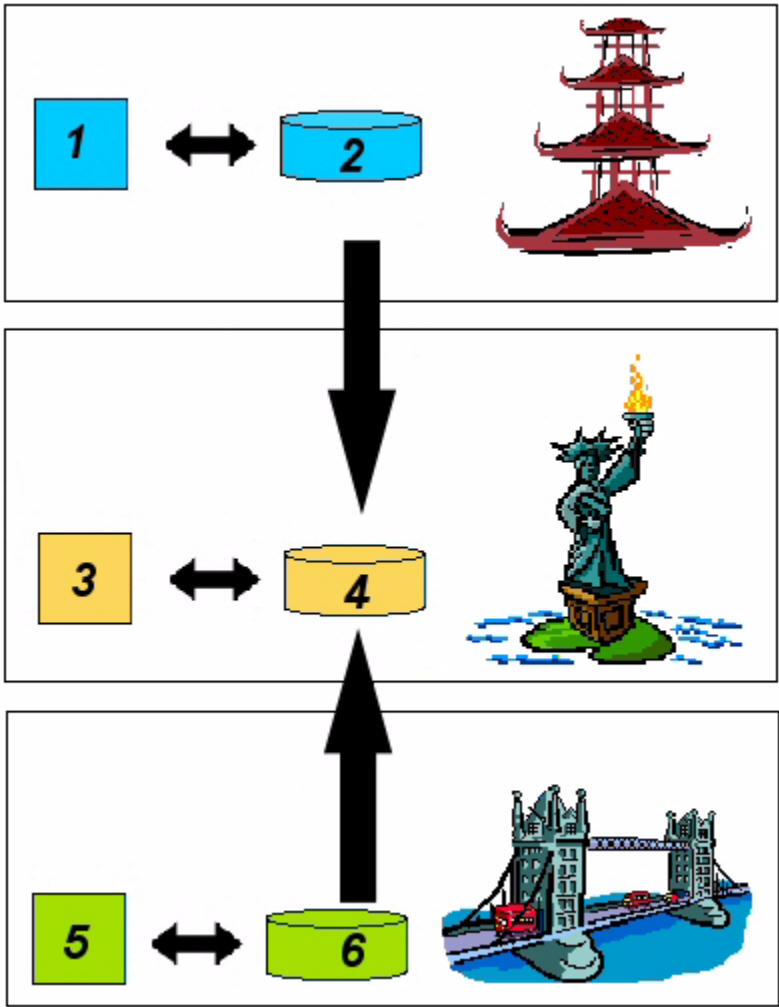


Figure 17 An Example of Global Reporting

Table 17 Description of “An Example of Global Reporting”


Item	Description
1	Management server for the SAN in Tokyo.
2	Database containing the local materialized views for the management server in Tokyo.
3	Management Server for the SAN in New York City.


Table 17 Description of “An Example of Global Reporting” (continued)

Item	Description
4	Database for the Management Server in New York City. It contains the local materialized views for New York City, in addition to the global views containing local data from Tokyo, London, and New York.
5	Management server for the SAN in London
6	Database containing the local materialized views for the Management Server in London.

Keep in mind the following when setting up global reporting:

- **Multiple Global Reporter Servers** - Global reporting can be set up at multiple sites. Your associates in London and Tokyo can also set up global reporting at their site so they obtain data gathered by management server in New York City.
- **Multiple Sites** - If you do have multiple Global Reporter servers, each Global Reporter server must contact each site that has the data the Global Reporter server wants to include in global reports. For example, assume you set up a Global Reporter server in Tokyo that collects data from a site in London. You also have another Global Reporter server in New York City. The Global Reporter server in New York City must contact the Global Reporter server in Tokyo and the site in London to obtain data from Tokyo and London. If the Global Reporter server in New York City contacts just the Global Reporter server in Tokyo, it does not obtain data from the site in London.
- **Data Gathering** - The Global Reporter server gathers data from sites as scheduled or when the **Refresh Now** button on the Global Reporter tab is clicked. The management server pulls data from a site's local materialized views, which are snapshots of relevant data from the site's local database. During a refresh, the Global Reporter server contacts each site and attempts to pull data from its materialized views. If a site's local view is refreshing, the Global Reporter server skips pulling the data from that view and proceeds pulling data from the remaining views. The Global Reporter server then attempts to pull data from the skipped views.
- **Firewalls** - If a site is behind a firewall, the port for the Oracle TNS listener must be open. For example, if the site uses port 1521 for its TNS listener, this port must be open on the firewall.
- **Contact Status** - The Contact Status column on the Global Reporter tab only verifies if the management server was able to contact the site. It does not specify if the refresh was successful. For example, if the Global Reporter server is able to contact the site, but the refresh fails. The Contact Status column still reports “SUCCESS” because the site was able to be contacted. If you have trouble contacting the site, try pinging it to verify the network connection is working. If ping works, verify the management server on the site is running.
- **Unable to Contact Site** - If the management server is unable to contact one of the sites in the Global Reporting list, the refresh process will not start. You can verify if a site can still be contacted by clicking the **Test** button for the site. You may want to try pinging the machine and

verify that management server on the remote server is running. If the site cannot be contacted, remove the obsolete site from the list by clicking the  button.


- **Security** - Users, whose role allows them access to view global reporting, can view all the elements throughout the enterprise. Grant access to viewing global reports only to those, who should be allowed to view all elements. Users belonging only to the following roles are given default privileges to view Global Reporter: CIO, Domain Admin, and Storage Administrator.
- **No data is displayed** - The Global Reporter server gathers information the site has collected. If the site has not collected any data, the Global Reporter server cannot obtain information from that site. Also, the Global Reporter server cannot display data until it has refreshed its view. The views are created at installation time and are empty until the view is refreshed. You can perform an immediate refresh by clicking the **Refresh Now** button on the Global Reporter server.
- **Data is outdated** - The Global Reporter server is as up to date as the sites. If the views on the sites are outdated, the Global Reporter server gathers this outdated information. If the management server is unable to contact the site, the Global Reporter server uses data from its last refresh.
- **Removing Sites** - To delete a site, click the corresponding  button for the site. Then, click the **Refresh Now** button near the bottom of the page. If you do not click the **Refresh Now** button and you attempt to create a new site with the same IP address as the deleted site, the management server tells you the site already exists. Also, when you delete a site, site names are not removed from the `tnsnames.ora` file. Names of deleted sites left in the `tnsnames.ora` file do not impact the user interface or the performance of the product. See ["Remote Sites Are Not Removed from the tnsnames.ora File"](#) on page 136 for more information.
- **Editing Sites** - You can edit the port number and site name of a site, but you cannot edit its IP address.

To set up global reporting:

1. The Global Reporter server contacts sites through the IP address/DNS name and Oracle TNS listener port. Gather the IP address and port number of the Oracle TNS listener port that each site uses.
2. On the Global Reporter server, go to the Global Reporter tab by clicking **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager.
3. Click the **Global Reporter** tab.
4. Click the **New Site** button.
5. Enter the information you gathered for a site in the first step.
 - **IP Address** - The IP address or DNS name of a server.
 - **Port number (optional)** - The Oracle TNS listener port the site uses. If this field is left blank, the management server assumes the database on the site uses port 1521.
 - **Site Name** - A name that includes the location of the site. For example, London1. Since data in the global reports is grouped by the site name, it is recommended you provide a unique site name to differentiate the sites.

Keep in mind the following:

- The remote site is not required to have global reporting enabled in its license.

- If you want data from the Global Reporter server included in global reports, add the local management server. Enter “localhost” as an IP Address/DNS name for your local management server.
6. Click **OK**.
 7. When you are done, click **OK**.
The management server verifies it can contact the site and it checks the version of the management server the site is running. The management server then adds the site to its list for global reporting.
 8. Repeat steps 4 and 5 for each site you want to add.
 9. Enable global reporting by selecting the **Enable** option on the Global Reporter tab.
 10. Set the time you want the refresh to start by doing the following:
 - a. Click the  button. In the **Time** field, type a time in the 24-hour format with hours and minutes separated by a colon. For example, 23:15, for 11:15 p.m.
 - b. Select a date you want the job to start.
The date is selected.
 - c. Click the **Set** button.
If you click the Set button after the time has passed, you must reset the time.
The time and date you selected are displayed in the **Next Schedule Run** field.
 11. Set up a repeat interval by typing a number in the **Repeat Interval** field and selecting a unit of measurement. For example, if you want the Global Reporter server to check the views of sites daily, you would type 1 in the **Repeat Interval** field and select **day(s)** as a unit.

IMPORTANT: You must select a repeat interval that is 12 hours or more.

12. When you are done with your changes, click the **Save Changes** button.
Your changes are saved in case the browser is closed.
13. You must modify the listener.ora file at each remote site, as described in the following steps. For example, assume you have three remote sites. You must log onto each of these remote sites and modify the listener.ora file at each remote site, as described in the following steps:
 - a. Log onto the remote site.
 - b. Stop the service for the management server running at the remote site.
 - c. Stop the listener service for Oracle:
Go to the Services window to stop the OracleOraHome92TNSListener service.
 - d. Open the following file in a text editor on the computer at the remote site:
`%ORA_HOME%\network\admin\listener.ora`
 - e. **After** `(ADDRESS = (PROTOCOL = TCP)(HOST = localhost)(PORT = 1521))`, add the following line:
`(ADDRESS = (PROTOCOL = TCP)(HOST = 192.168.10.1)(PORT = 1521))`
where 192.168.10.1 is the IP address of the local host server. Replace 192.168.10.1 with the IP address of your local host.

The text should now appear as follows:

```
LISTENER =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS_LIST =
        (ADDRESS = (PROTOCOL = TCP) (HOST = localhost) (PORT = 1521))
        (ADDRESS = (PROTOCOL = TCP) (HOST = 192.168.10.1) (PORT = 1521))
      )
    )
  )
```


f. Save the file and exit.

g. Start the listener service for Oracle:

Go to the Services window to start the OracleOraHome92TNSListener service.

h. Start AppStorManager.


14. To obtain information immediately, click the **Refresh Now** button (**Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager).

15. When the Global Reporter server is done with updating its views, you can view the global reports. To view global reports, click **Reporting** (). Expand the Global node in the tree to view the reports.

You can filter the data viewed in the reports. See "[Filtering Data in Global Reports](#)" on page 342.

Editing Remote Site Information

To modify the information for your remote servers:

1. Click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager.
2. Click the **Global Reporter** tab.
3. Click  corresponding to the remote site you want to modify.
4. Modify the following information for your remote servers that are running the management server:
 - **IP Address** - The IP address or DNS name of a server. If you change the IP address, you must modify the `listener.ora` file on the remote server, as described in the following step.
 - **Port number (optional)** - The Oracle TNS listener port the remote server uses. If this field is left blank, the management server assumes the database on the remote server uses port 1521.
 - **Site Name** - A name that includes the location of the site. For example, London1. Since data in the enterprise reports is grouped by the site name, it is recommended you provide a unique site name to differentiate the sites.
5. Click **OK** when done.

6. If you changed the IP address of the remote server, you must edit the `listener.ora` file on that remote server as described in the following steps:
 - a. Log onto the remote site.
 - b. Stop the service for the management server running at the remote site.
 - c. Stop the listener service for Oracle (OracleOraHome92TNSListener) at the remote site.
 - d. Stop OracleServiceApplQ at the remote site.
 - e. Open the following file in a text editor on the computer at the remote site:
`%ORA_HOME%\network\admin\listener.ora`
 - f. Edit the IP address as shown in the following line.


```
(ADDRESS = (PROTOCOL = TCP) (HOST = 192.168.10.1) (PORT = 1521))
```
 - g. Save the file and exit.
 - h. Start the listener service for Oracle (OracleOraHome92TNSListener).
 - i. Start AppStorManager.

Remote Sites Are Not Removed from the tnsnames.ora File

Each time you add a site for global reporting, its contact information is added to the `tnsnames.ora` file. When you delete a site, site names are not removed from the `tnsnames.ora` file. Names of deleted sites left in the `tnsnames.ora` file do not impact the user interface or the performance of the product. For example, assume you added the site “remotesiteA” for global reporting. It would appear in the `tnsnames.ora` file as the following:

```
remotesiteA =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = 192.168.1.2) (PORT = 1521))
    )
    (CONNECT_DATA =
      (SID = applq)
    )
  )
```

Now, let's assume you removed “remotesiteA” from the user interface by clicking the

corresponding  button for the site and then the **Refresh Now** button near the bottom of the Global Reporter page (**Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager). Even though the remote site has been removed from the user interface, it still appears in the `tnsnames.ora` file. If you add “remotesiteA” as a remote site for global reporting again, another listing will be added to the `tnsnames.ora` file. These dual listings do not negatively impact the product.

Managing Performance Collection

This section describes the following:

- “[Managing Performance Collectors](#)” on page 137
- “[Starting Performance Collectors](#)” on page 138

- ["Stopping Performance Collectors"](#) on page 139
- ["Setting the Date and Time for Performance Collectors"](#) on page 139
- ["Viewing Data Aging Statistics for Performance"](#) on page 140

Managing Performance Collectors

The management server uses performance collectors to gather information for Capacity Manager and Performance Manager charts, as well as for monitoring. Use Data Collection page for Performance collectors to stop and start collectors as well as schedule when they run. Each row in the table corresponds to a collector. The Element Type column displays which element the collector gathers data for, and the Statistics column displays which statistics for that element the collector is tasked to gather.

For example, assume you see host_A in the Element Type column and Processor Utilization in the Statistics column. That collector gathers data about processor utilization for host_A.

You can have multiple collectors for an element. Let's go back to the example of host_A. Assume this host has another collector. Underneath the first row there is another row for the second collector. The Element Type is host_A, but this time the following statistics are listed: Physical Memory Used, Free Physical Memory, Virtual Memory Used, Free Virtual Memory. The second collector gathers information about all of those statistics listed for host_A.

You can see the result of these statistics in Capacity Manager and Performance Manager. When you select one of the following options from the Period combo-box, you can view the information these collectors gathered. Real time statistics are not affected by these collectors


- **Last Hour***
- **Last 24 Hours***
- **Last 7 Days**
- **Last Month**
- **Last Year**

*Only available in Performance Manager

IMPORTANT: All collectors are stopped during "Discovery Data Collection." This means that during "Discovery Data Collection" data, such as for Performance Manager, is not updated. Historical collectors, such as those available from the Configuration tab, are restarted when they are stopped during "Discovery Data Collection. Any charts that were active in Performance Manager when "Discovery Data Collection" was started are not restarted.


To manage performance collectors, click **Optimize > Storage Essentials > Performance Data Collection**.

Table 18 About Performance Collectors

Column Heading	Description
Element	The name of the element from which this collector gathers information.
Element Type	The type of element from which the collector gathers information.
Statistics	The statistics that the collector is responsible for providing information.
Next Scheduled Run	Displays the date and time when the collector is scheduled to run.
Interval	The interval in minutes the collector runs.
Running?	Gives the status of the collector. Collectors that are running display a check mark in this column.
Edit	<p>To edit the schedule for running a collector, click the Edit  button. Then set the date and time. See "Setting the Date and Time for Performance Collectors" on page 139 for more information.</p> <p>Important: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the first day of the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.</p>
Action	<p>Displays one of the following buttons:</p> <ul style="list-style-type: none"> • Stop - Stops the collector. The corresponding reports display information only gathered previously. See "Stopping Performance Collectors" on page 139. • Start - Starts the collector. When you start a collector, it begins gathering information for its corresponding reports. See "Starting Performance Collectors" on page 138.

Starting Performance Collectors

To start a collector:

1. Access the page for performance collectors page (**Optimize > Storage Essentials > Performance Data Collection**).
2. Click the **Start** button corresponding to the collector you want to start.
3. Click the calendar icon, .
4. Set the date and time as described in the topic, "[Setting the Date and Time for Performance Collectors](#)" on page 139.

5. Type a repeat interval and then select a unit of measurement from the drop-down menu.
The repeat interval determines how often the collectors gather the data.
6. If you are asked to provide a proxy host, do the following:
 - a. Select a proxy host by clicking the **Browse** button.
 - b. Select a proxy host from the drop-down menu and then click **OK**.
 - c. Click **OK** again to set the time for starting the collector.
 - d. If you do not see any hosts displayed verify you have the latest CIM Extension version installed and running on a host that can access the Engenio storage system.
7. Click **OK**.
8. To start more than one collector at once, select more than one collector on the **Data Collection** tab and then click the **Start Collectors** button.

Stopping Performance Collectors

When you stop a collector, the management server stops gathering information for which the collector is responsible. For example, if a performance collector is not running, its corresponding statistics are no longer receiving information to display. One of the following occurs:

- If there was originally no information gathered for the statistic, no data appears for that statistic.
- If information was previously gathered for the report, old data appears for that statistic.



To stop a collector:

1. Access the page for performance collectors (**Optimize > Storage Essentials > Performance Data Collection**).
2. Click the **Stop** button corresponding to the collector you want to stop.
The collector stops gathering information for its corresponding reports.

To stop more than one collector at once, select more than one collector and then click the **Stop Collectors** button.

Setting the Date and Time for Performance Collectors

To set the date and time, do the following:




1. Click **Optimize > Storage Essentials > Performance Data Collection**.
2. Click the **Edit** button, .
3. Click the calendar icon, .
4. Enter the time in the time field. Make sure the time resembles a 24-hour clock, for example 22:00 for 10 p.m.
5. Click the date.

The date is highlighted in pink.

You can navigate the calendar as follows:



- Displays the same month in the previous year

-  Displays the previous month
-  Displays the next month
-  Displays the same month in the following year

6. When you are done, click the **Set** button.
The management server displays the date and time.

IMPORTANT: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.



Figure 18 Selecting the Date and Time from the Calendar

In the previous figure, the calendar sets the date and time at 1:56 p.m. February 13, 2003.

Viewing Data Aging Statistics for Performance

Data Aging includes Data Rollup and Garbage Collection. Data Rollup controls how often a set of data is summarized. For example, hourly data is rolled into the daily table periodically. Garbage Collection refers to how long a set of data is preserved before it is permanently removed from the database.


The settings on this page control data aging for both reports and performance statistics.

IMPORTANT: Do not modify the data on this page unless instructed by customer support. Changing them incorrectly can adversely affect the management server.

To view data aging statistics:

1. Click **Optimize > Storage Essentials > Performance Data Collection**.
2. Click the **Data Aging** tab at the top of the screen.

IMPORTANT: Perform the following steps **ONLY** if customer support has instructed you to modify one of the collectors on the page:

3. Click the **Edit**  button.
4. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Click the date on which you want to run the collector to run the next time. Today's date is highlighted in pink. Click **Set**.
The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd format.
If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.
5. Click the **Enable** option.

NOTE: If you are not allowed to disable the collector, the Enable option is unavailable.

6. To change the repeat interval, In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - **Second(s)**
 - **Minute(s)**
 - **Hour(s)**
 - **Day(s)**
 - **Week(s)**
7. (Garbage Collection only) To change how long the data is preserved, type an interval in the **Preserve** field, and then select one of the following from the drop-down menu to the right of the field:
 - **Second(s)**
 - **Minute(s)**
 - **Hour(s)**
 - **Day(s)**
 - **Week(s)**
8. Click the **OK** button.

7 Database Maintenance and Management

This chapter contains the following topics:

- ["Database Maintenance Window"](#) on page 143
- ["Overview of Backups"](#) on page 144
- ["Database Mode"](#) on page 145
- ["Architectural Overview of RMAN Backups"](#) on page 146
- ["Performing an RMAN Hot Backup"](#) on page 148
- ["Scheduling RMAN Hot Backups of the Database"](#) on page 148
- ["Viewing Results from RMAN Backup"](#) on page 149
- ["About the Database Admin Utility"](#) on page 150
- ["Database Password"](#) on page 158
- ["Troubleshooting Listener and Database Connection Problems"](#) on page 159

Database Maintenance Window

Schedule a maintenance window of two to four hours weekly during off peak hours of operation to do the following database maintenance operation:

1. Stop the service for the management server. See ["Stopping the Service for the Management Server"](#) on page 11.
2. Access the Database Admin Utility. See ["Accessing the Database Admin Utility"](#) on page 150 for more information.
3. Using the Database Admin Utility, verify the database is in an open state and the listener is running. See ["Checking the Database and Listener Status"](#) on page 151.

If the database state is not OPEN and shows error get the following logs and then contact technical support:

- Log file for the Database Admin Utility (%MGR_DIST%/Tools/dbAdmin.log)
 - Database alert log which can be found in \oracle\admin\APPIQ\bdump, where ORACLE_BASE is c:\oracle
4. Reset the temporary tablespace, as described in ["Resetting the Temp and Undo Tablespace"](#) on page 153.
 5. Export the database, as described in ["Exporting the Database"](#) on page 152.
 6. If the database is running in archive mode, set the database to no-archive mode, as described in ["Setting the Archive Mode"](#) on page 154. Then, clean the \oracle\oradata\APPIQ\archive.

NOTE: The archive directory (\oracle\oradata\APPIQ\archive) only exists if you have previously set the management server to archive mode.

- 7. Return the database to archive mode, as described in “[Setting the Archive Mode](#)” on page 154.
- 8. Start the service for the management server.
- 9. If database is in archive mode, take a current RMAN backup by clicking the **Backup Now** button on the **Options > Storage Essentials > Manage Product Health > Disk Space > RMAN Backup** in HP Systems Insight Manager. See “[Performing an RMAN Hot Backup](#)” on page 148 and “[Architectural Overview of RMAN Backups](#)” on page 146. On successful completion of RMAN backup, the backup is saved to the following directory:
%ORA_HOME%\rman\current
- 10. Clean the following folders:
%ORA_HOME%\rman\backup1
%ORA_HOME%\rman\backup2

Overview of Backups

The management server provides the following backups:

IMPORTANT: Export and RMAN backups should be done regularly in combination.

Table 19 Description of Backups

Backup Type	Description	Files Backed Up	Database Mode
Export backup	Done through the Database Admin Utility. See “ Exporting the Database ” on page 152 for more information about exporting the database.	Database Schema, Oracle network Configuration files (tnsnames.ora, listener.ora), CIM repository, File System Viewer	Does not matter
RMAN HOT backup	The backup is referred to as being “hot” because the management server is still running while the backup is occurring. You can configure the RMAN backup to run by default. See “ Scheduling RMAN Hot Backups of the Database ” on page 148.	Database Files, Control files, Redo files, Archive Files, Oracle network Configuration files (tnsnames.ora, listener.ora), CIM Repository, File System Viewer	To do an RMAN hot backup, the management server must be set to archive mode. See “ Setting the Archive Mode ” on page 154.

Table 19 Description of Backups (continued)

Backup Type	Description	Files Backed Up	Database Mode
RMAN Cold backup	Done through the Database Admin Utility.	Same files as an RMAN HOT backup.	If the management server is set to no-archive mode, users can perform an RMAN cold backup. See "Setting the Archive Mode" on page 154.

Backup Destination and Operation for RMAN: The management server has three backup points available for RMAN backups. RMAN backup destinations are:

```
%ORA_HOME%\rman\current
%ORA_HOME%\rman\backup1
%ORA_HOME%\rman\backup2
```

The schedule backup writes in backup1 and backup2 folder in rotation and save now backup from SA keeps overwriting in current folder. See ["Architectural Overview of RMAN Backups"](#) on page 146 for more information about RMAN backups.

Database Mode

To facilitate the HOT RMAN Backup, you must change the database mode to archive mode, as described in ["Setting the Archive Mode"](#) on page 154. The default database archive destination is \oracle\oradata\APPIQ\archive. This destination can be modified as described in ["Changing the Archive Destination"](#) on page 157. Depending on the input/output of the data, archiving can be in the range of 2.5 GB to 10 GB per day. The archive folder gets cleaned during on a scheduled RMAN backup.

You can change the database mode to run in no archive mode with RMAN backup disabled, as described in ["Setting the Archive Mode"](#) on page 154. When you change the database to no archive mode, you reset the logs SCN, set the archive parameter in the database parameter file and disable the RMAN backup scheduler.

IMPORTANT: Export the database after switching to no archive mode. See ["Exporting the Database"](#) on page 152 for more information.

Keep in mind the following implications if you do decide to change the database to no archive mode.

- If you set the management server to **Disable Database Archive Mode and RMAN Backup**, it is up to you to back up the management server manually. If you forget to back up your management server and your management server fails, you will not have a database to import. To learn more about manual backups, see ["Running a Cold Backup"](#) on page 156. Export the database after you change the database mode to no archive. See ["Exporting the Database"](#) on page 152.

- Scheduled RMAN backup jobs do not run. If you change the database mode to no archive during an RMAN backup, the RMAN backup will error out.
- If the database fails as a result of a corrupt data file, the database can only be restored to the last export backup available. This requires recreating the database along with the import.

When you change the database to archive mode, you reset the logs SCN, set the archive parameter in the database parameter file and enable the RMAN backup scheduler. Take a current RMAN backup after switching to archive mode, as described in [“Performing an RMAN Hot Backup”](#) on page 148.

Architectural Overview of RMAN Backups

By default the management server does not backup the database automatically. If you enabled the database archive mode and RMAN backup as described in [“Setting the Archive Mode”](#) on page 154, the management server backs up the Oracle instance for the management server every three days and saves the backup for two weeks. The management server alternates the directories in which it saves the backup. When the management server first performs a scheduled backup, the backup is saved in %ORA_HOME%\rman\backup1. The next time the management server performs a scheduled backup, it is saved %ORA_HOME%\rman\backup2. The management server saves the backup in alternating directories (backup1 and backup2), so you always have a copy of the last backup and the previous backup. To learn how to change the frequency of the scheduled backups, see [“Scheduling RMAN Hot Backups of the Database”](#) on page 148.

You can back up the database at any time by clicking the **Backup Now** button on the Database tab. When you back up the database using this technique, the backup is saved only in %ORA_HOME%\rman\current. To recover the database, contact customer support.

Table 20 Backup Directories in %ORA_HOME%\rman

Directory	Contains
current	The backup when the Database Server Backup button was clicked.
log	A log of when the backup was done.
backup1	Information from the automatic backup (alternating day).
backup2	Information from the automatic backup (alternating day).

Let's assume you recently installed the management server and you have not done any backups. You have schedule the backups to take place every three days. You performed a backup, and it is stored in the backup1 folder. The next scheduled backup occurs on day 4 and it is saved in the backup2 folder. If your database fails, you can restore the database from day 1 or day 4. Let's assume you have a scheduled backup on day 7, which is saved to the backup1 folder. This backup replaces the backup from day 1. The available backups are from day 4 and 7. Assume

you click the **Backup Now** button on day 8, the backup is saved in the **CURRENT** folder because the backup is recording the current state of the database. If your database fails, you can restore the database from day 4, 7 or 8, as described in the following table:

Table 21 Backup Example

Day	Backup Type	Backup 1	Backup 2	Current	Available Backup
Day 1	Scheduled	Day 1 backup	-----	-----	Day 1 Backup
Day 4	Scheduled	-----	Day 4 backup	-----	Backups from Days 1 and 4
Day 7	Scheduled	Day 7 backup	-----	-----	Backups from Days 4 and 7
Day 8	Save Now	-----	-----	Day 8	Backups from Days 4, 7, and 8
Day 10	Scheduled	-----	Day 10 backup	-----	Backups from Days 7, 8, and 10

Keep in mind the following:

- Only one user at a time can back up the database.
- The management server archives files for the backup in a separate directory. Do not modify the files in this directory (`\oracle\oradata\APPIQ\archive`). For average database activity, the management server requires at least 100 GB of disk space for archive files. If there is higher database activity than average, more disk space may be required.

The following is saved during the backup:

- **Management server RMAN backup files** - These files contain information about the elements your management server monitors.
- **Oracle Network Configuration Files** -The configuration files are `tnsnames.ora` and `listener.ora`.
- **CIM Repository**
- **Property files, such as `jboss.properties`**
- **`rmanbackup.log` file**
- **`spfileappiq.ora` file**

To back up the database manually:

1. Before you can back up the database manually, you must enable database archive mode and RMAN backup as described in "[Setting the Archive Mode](#)" on page 154.

2. In the management server, click **Options > Storage Essentials > Manage Product Health** in HP Systems Insight Manager.
3. Select **RMAN Backup** in the Product Health tree.
4. Click the **Backup Now** button.
The database is backed up.

Performing an RMAN Hot Backup

You can perform an RMAN hot backup instantly. The backup is referred to as being “hot” because the management server is still running. When you perform an RMAN hot backup, the following files are backed up: database files, control files, redo files, archive files, Oracle network configuration files (tnsnames.ora, listener.ora), CIM repository, and File System Viewer.

NOTE: The buttons on the RMAN backup page appear disabled if the database archive mode is disabled. See [“Setting the Archive Mode”](#) on page 154 for more information about changing the archive mode.

To perform an RMAN Hot Backup:

1. Verified you have enabled database archive mode and RMAN backup as described in [“Setting the Archive Mode”](#) on page 154.
2. Click **Options > Storage Essentials > Manage Product Health** in HP Systems Insight Manager.
3. Select **RMAN Backup** in the Product Health tree.
4. Click **Backup Now**.

The management server performs an RMAN backup.

When the backup is complete, the Setup tab is refreshed with the status of the manual hot backup. The Results tab is also updated with the status of the RMAN hot backup. The Results tab displays the status of the previous RMAN hot backups (manual and scheduled).

When the management server performs a hot RMAN backup,

Scheduling RMAN Hot Backups of the Database

The management server lets you schedule an RMAN hot back up of the database, as described in the following steps. The backup is referred to as being “hot” because the management server is still running.

IMPORTANT: Before you can schedule the RMAN backup, you must enabled database archive mode and RMAN backup as described in [“Setting the Archive Mode”](#) on page 154. The buttons on the RMAN backup page appear disabled if the database archive mode is disabled.

To learn more about backing up the database, see [“About the Database Admin Utility”](#) on page 150.

1. Verified you have enabled database archive mode and RMAN backup as described in ["Setting the Archive Mode"](#) on page 154.
2. Click **Options** > **Storage Essentials** > **Manage Product Health** in HP Systems Insight Manager.
3. Select **RMAN Backup** in the Product Health tree.

4. Click the calendar icon, .

5. In the **Time** field, type the time in 24-hour format with the hour and minutes separate by a colon, for example, 22:15. Click the date on which you want to run the next backup of the database. Today's date is highlighted in pink. Click **Set**.

The date and time appear in the **Next Scheduled Run** field in the yyyy-mm-dd format.

If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.

6. In the **Repeat Interval** field, type an interval. Select one of the following units from the drop-down menu:
 - **Second(s)**
 - **Minute(s)**
 - **Hour(s)**
 - **Day(s)**
 - **Week(s)**

NOTE: The minimal interval you can schedule is one day. If you select **Hour(s)**, **Minute(s)** or **Second(s)**, you must enter an interval that equals more than a day. For example, if you select **Hour(s)**, you must enter 24 or more. Just as if you select **Minute(s)**, you must enter 1440.

7. Click the **Enable** option.
8. Click the **Save Schedule** button.

You can always disable the schedule by clearing the **Enable** option.

When the scheduled RMAN hot backup is complete, the Results tab is updated with the status of the RMAN hot backup. The Results tab displays the status of the previous RMAN hot backups (manual and scheduled).

Viewing Results from RMAN Backup

You can determine the results of the RMAN backups. The Results tab for RMAN backup provides the data and time of the RMAN backup, in addition to its status and backup folder.

To view the results from RMAN backup:

1. Click **Options** > **Storage Essentials** > **Manage Product Health** in HP Systems Insight Manager.
2. Select **RMAN Backup** in the Product Health tree.

3. Click the **Results** tab in the RMAN Backup window.

The following information is displayed:

- Data/Time of the backup
- Status of the backup
- Backup folder

About the Database Admin Utility

The Database Admin Utility lets you easily manage your database from restoring it from a cold backup to resetting the temp tablespace. This tool provides flexible importing and exporting features, which let you save time. The Database Admin Utility requires you to stop the AppStorManager service so that it can manage the database optimally.

To learn more about the Database Admin Utility, see the following topics:

- ["Accessing the Database Admin Utility"](#) on page 150
- ["Refreshing the Database Admin Utility"](#) on page 150
- ["Checking the Database and Listener Status"](#) on page 151
- ["Changing the Password of System Accounts"](#) on page 151
- ["Exporting the Database"](#) on page 152
- ["Importing the Database"](#) on page 152
- ["Reinitializing the Database"](#) on page 153
- ["Resetting the Temp and Undo Tablespace"](#) on page 153
- ["Restarting the Database"](#) on page 154
- ["Restoring a Cold Backup"](#) on page 154
- ["Setting the Archive Mode"](#) on page 154
- ["RMAN Restore"](#) on page 155
- ["Running a Cold Backup"](#) on page 156
- ["Changing the Archive Destination"](#) on page 157
- ["Viewing Logs from the Database Admin Utility"](#) on page 157
- ["Resetting the Log"](#) on page 157
- ["Warning Messages During Reinitializing the Database"](#) on page 158

Accessing the Database Admin Utility

To access the Database Admin Utility:

1. Stop the AppStorManager service.
2. To access the database utility on Windows, go to the [Install_Dir]\Tools directory on the management server and double-click dbAdmin.bat, where [Install_Dir] is the directory into which you installed the management server.

Refreshing the Database Admin Utility

The Database Admin Utility requires the following are not in use:

- JBoss
- CIMOM

If any of these components are in use, the Database Admin Utility will not work. If you are shown an error message when you start the Database Admin Utility, stop the AppStorManager service. Then, click the **Refresh** button.

Checking the Database and Listener Status

To find the database and listener status:

1. Access the Database Admin Utility as described in "[Accessing the Database Admin Utility](#)" on page 150.
2. Click **Check Database Status** in the left pane.
3. Type the SYS password in the **Password for User SYS** field.
4. Click **Check Database Status**.
The database and listener status is shown in the lower pane.
5. To clear the pane, click **Clear All**.

Changing the Password of System Accounts

The management server uses the following accounts to access and manage the database for the management server. You should change the passwords to these accounts to prevent unauthorized access.

- **SYS** - Used for the management server database creation. Default password: `change_on_install`
- **SYSTEM** - Used for management server database creation, in addition to database import, export and re-initialization. Default password: `manager`
- **RMAN_USER** - Used for RMAN backup and restore. This user has sys privilege. Default password: `backup`
- **DB_SYSTEM_USER** - Used for all the database activity, including establishing a connection to the management server database. Default password: `password`

You must change the passwords of the SYS, SYSTEM, RMAN_USER, and DB_SYSTEM_USER accounts by using the Database Admin Utility tool, so the management server is aware of the changes. The passwords must also follow the following guidelines. Do not change the password for one of these accounts by using Oracle. Make sure you keep the new passwords in a safe location, as it is your responsibility to remember the Oracle passwords.

The management server requires the password to have the following characteristics:

- a minimum of three characters
- starts with a letter
- contains only letters, numbers and underscores (`_`)
- does not start or end with an underscore (`_`)

To make the management server aware of the new password:

1. Access the Database Admin Utility as described in "[Accessing the Database Admin Utility](#)" on page 150.
2. Click **Change Passwords** in the left pane.
3. Select an account name from the **User Name** combo box.
4. Type the current password in the **Old Password** field.
5. Type the new password in the **New Password** field.
6. Retype the password in the **Confirm Password** field.
7. Click **Change**.

The Database Admin Utility changes the password for the specified account.

Exporting the Database

Use the Database Admin Utility to save the database in a format that can be imported.

IMPORTANT: The export database feature should not be used for backing up the server. If you are interested in backing up the database for disaster recovery, use the Database Server Backup feature. See "[About the Database Admin Utility](#)" on page 150 for more information.

To export the database:

1. Access the Database Admin Utility as described in "[Accessing the Database Admin Utility](#)" on page 150.
2. Click **Export Database** in the left pane.
3. Click the **Browse** button to select a file path. Then, type a file name in the **File name** field. Then, click the **Open** button.

IMPORTANT: Select a directory outside of the directory tree of the management server. This way if you remove the management server, you will not lose the saved database.

The file name and its path is displayed in the Database Admin Utility. The .zip file extension is automatically added to the file name.

4. Type the password of the SYSTEM account. The default password of the SYSTEM account is the following: `manager`
You are notified when the database process is complete.
5. Select **Clear Report Cache** if you do not want the report cache to be included with the database you are exporting. When a user imports this database, their report cache will be empty until the report cache is refreshed (**Reports > Storage Essentials > Report Configuration > Report Cache**). This option may save you time with exporting the database if your database includes a large amount of data for reports.
6. Click **Export Database**.

Importing the Database

You can revert to an earlier configuration by uploading a file (*.zip) containing the database information.

The software stores a snapshot of the data in its database. Since this file is a snapshot of the network at a certain time, it may not contain your most current network configuration. If you want to view an up-to-date network configuration and the latest information about the elements, perform discovery and Discovery Data Collection.

To import a database:

1. Access the Database Admin Utility as described in "[Accessing the Database Admin Utility](#)" on page 150.
2. Click **Import Database** in the left pane.
3. Click the **Browse** button. Select a database file with a .zip extension to import. Then, click the **Open** button.

The file name is displayed in the Database Admin Utility.

4. Type the password of the SYSTEM account. Then, click **OK**. The default password of the SYSTEM account is the following: `manager`

You are notified when the database process is complete.

5. Select **Populate Report Cache** if you want to refresh reports cache during import. Keep in mind that the amount of time to import the database may increase if you select this option and the database has a large amount of data for reports. Alternatively you can refresh the report cache from the management server (**Reports > Storage Essentials > Report Configuration > Report Cache**).
6. Click the **Import Database** button.

Reinitializing the Database

Keep in mind the following:

- Reinitializing the database removes everything from the database. This is not recommended unless you are sure about what you are doing. It is strongly suggested you export the database before you reinitialize it. See "[Exporting the Database](#)" on page 152 for more information on how to save the database.
- When you reinitialize the database, all users are logged out of the management server.
- Ignore the warning messages in the command prompt window that pops up when the `dbAdmin.bat` file is double-clicked. See "[Warning Messages During Reinitializing the Database](#)" on page 158 for more information.

To clear the database:

1. Access the Database Admin Utility as described in "[Accessing the Database Admin Utility](#)" on page 150.
2. Click **Re-initialize Database** in the left pane.
3. Type the password of the SYSTEM account. Then, click **OK**. The default password of the SYSTEM account is the following: `manager`

You are notified when the re-initialization is complete.

4. Click the **Re-initialize Database** button.

Resetting the Temp and Undo Tablespace

The temporary and undo tablespace may grow large due to high database activity. You should regularly reset the temp and undo tablespace to its initial value, as described in this section.

To reset the tablespace:

1. Access the Database Admin Utility as described in "[Accessing the Database Admin Utility](#)" on page 150.
2. Click **Reset Temp Tablespace** in the left pane.
3. Type the SYS password in the **SYS Password** field.
4. Click the **Reset Temp Tablespace** button.

Restarting the Database

You may sometimes need to restart the ApplQ instance of the database. Use this feature in the Database Admin Utility when the database is down or you need to shutdown and then start the database.

To restart the database:

1. Access the Database Admin Utility as described in "[Accessing the Database Admin Utility](#)" on page 150.
2. Click **Restart Database Server** in the left pane.
3. Type the SYS password in the **SYS Password** field.
4. Click the **Restart Database Server** button.

Restoring a Cold Backup

If you performed an RMAN cold backup, follow the steps in this section to restore the RMAN cold backup. You can only perform a cold backup if you selected **Disable Database Archive Mode and RMAN Backup**. See "[Setting the Archive Mode](#)" on page 154 for information about changing the archive mode. The backup is referred to as being "cold" because the management server is not running while the backup is occurring.

To restore a cold backup:

1. Access the Database Admin Utility as described in "[Accessing the Database Admin Utility](#)" on page 150.
2. Click **Restore Cold Backup** in the left pane.
3. Click the **Browse** button.
4. In the **File Name** field, provide a directory path containing the cold backup (may automatically be populated in some Web browsers), but do not provide a file name. In this release the cold backup is saved in the COLDBACKUP directory in the path specified by the person who did the cold backup. The RMAN backup consists of several files. The Database Admin Utility automatically notices the backup files in the directory provided.

5. Type the password of the SYS account.
6. Click **Run Cold Backup**.

Setting the Archive Mode

By default the management server database runs in no-archive mode, which requires you to backup the database manually by using a cold RMAN backup. A cold RMAN backup is an RMAN backup without the management server running. If you want to take an RMAN hot backup of the database, switch the database to archive mode. An RMAN hot backup is with the database running while the backup is occurring.

Keep in mind the following implications if you decide to leave the database in no archive mode.

- If you leave the database in no archive mode (**Disable Database Archive Mode and RMAN Backup**), it is up to you to back up the management server manually. If you forget to back up your management server and your management server fails, you will not have a database to import. To learn more about manual backups, see ["Running a Cold Backup"](#) on page 156. Export the database after you change the database mode to no archive. See ["Exporting the Database"](#) on page 152.
- Scheduled RMAN backup jobs do not run. If you change the database mode to no archive during an RMAN backup, the RMAN backup will error out.
- If the database fails as a result of a corrupt data file, the database can only be restored to the last export backup available. This requires recreating the database along with the import.

To change the archive mode:

1. Access the Database Admin Utility as described in ["Accessing the Database Admin Utility"](#) on page 150.
2. Click **Archive Mode Switch** in the left pane.
3. Select one of the following:
 - **Enable Database Archive Mode and RMAN Backup** - Select this option if you plan to run database in archive mode and automated backup while the management server is running. If database is already in archive mode, you can use this option to clean up archive files. See ["About the Database Admin Utility"](#) on page 150 for more information about the automatic backups. When you change the database to archive mode, you reset the logs SCN, set the archive parameter in the database parameter file and enable the RMAN backup scheduler. Take a current RMAN backup after switching to archive mode (**Options > Storage Essentials > Manage Product Health > Disk Space > RMAN Backup > Backup Now** in HP Systems Insight Manager) as described in ["Performing an RMAN Hot Backup"](#) on page 148.
 - **Disable Database Archive Mode and RMAN Backup** - Select this option if and only if you always shut down the management server prior to a backup. With database running in no-archive mode, you can only cold back up your database. When you change the database to no archive mode, you reset the logs SCN, set the archive parameter in the database parameter file and disable the RMAN backup scheduler. See ["Running a Cold Backup"](#) on page 156 for more information. Export the database after switching to no archive mode. See ["Exporting the Database"](#) on page 152.
4. Type the SYS password in the **SYS Password** field.

5. Click the **Change Settings** button.

RMAN Restore

The Database Admin Utility lets you restore the management server database from a previously scheduled (hot) RMAN backup, which are stored in the following directories by default: backup1, backup2, and current. You may not have a previously scheduled (hot) RMAN backup if you selected **Disable Database Archive Mode and RMAN Backup**.

To restore the database:

1. Access the Database Admin Utility as described in "[Accessing the Database Admin Utility](#)" on page 150.
2. Click **RMAN Restore** in the left pane.
3. Select one of the following directories to restore:
 - **Current** - The last restore from when the **Database Server Backup** button on the Database tab was clicked.
 - **Backup1***
 - **Backup2***

*Information from the automatic backup (alternating weeks). See "[About the Database Admin Utility](#)" on page 150 for more information about the automatic backup.
4. Type the SYS password in the **SYS Password** field.
5. Click **Restore RMAN**.

The Database Admin Utility restores the selected database.

Running a Cold Backup

If you are running the database in no-archival mode, you should perform a cold backup frequently. The backup is referred to as being "cold" because the management server is not running while the backup is occurring.

NOTE: This feature does not run if the **Enable Database Archive Mode and RMAN Backup** option is selected. The **Enable Database Archive Mode and RMAN Backup** option performs an RMAN backup periodically. See "[About the Database Admin Utility](#)" on page 150 and "[Setting the Archive Mode](#)" on page 154 for more information.

The following is saved during a cold RMAN backup:

- **Management server RMAN backup files** - These files contain information about the elements your management server monitors.
- **Oracle Network Configuration Files** -The configuration files are `tnsnames.ora` and `listener.ora`.
- **CIM Repository**
- **Property files, such as jboss.properties**
- **rmanbackup.log file**

- **spfileappiq.ora** file

To run a cold backup:

1. Access the Database Admin Utility as described in "[Accessing the Database Admin Utility](#)" on page 150.
2. Click **Run Cold Backup** in the left pane.
3. Click the **Browse** button to select a file path.
4. In the **File Name** field, provide a directory path (may automatically be populated in some Web browsers), but do not provide a file name.

IMPORTANT: The management server saves the backup in a directory called COLDbackup in the path you specified. Any pre-existing content in this directory, such as previous cold backups, is removed.

5. Type the password of the SYS account.
6. Click **Run Cold Backup**.

Changing the Archive Destination

The current archive directory by default is `\oracle\oradata\APPIQ\archive`. Over time your database will grow. If you feel you are running out of space, you can add a new volume and change the archive destination to a new volume, as described in the following steps:

1. Access the Database Admin Utility as described in "[Accessing the Database Admin Utility](#)" on page 150.
2. Click **Set Archive Destination** in the left pane.
3. Click the **Browse** button to select a file path.
4. Type the password of the SYS account.
5. Click **Set Archive Destination**.

Viewing Logs from the Database Admin Utility

You can determine the status of a process executed by the Database Admin Utility by looking at its logs. To view the log from the Database Admin Utility, click **View Log** in the Database Admin Utility.

1. Access the Database Admin Utility as described in "[Accessing the Database Admin Utility](#)" on page 150.
2. Click **View Log**.
A separate window appears displaying the logs.
3. You can refresh the logs by clicking the **Refresh** button in the separate window displaying the logs. To clear this window, click **Reset Log** in the Database Admin Utility window.

Resetting the Log

When you click the **View Logs** button, a separate window appears displaying the logs. You can clear this separate window by clicking **Reset Log**. When you are asked if you want to reset the log, click **Yes**.

Warning Messages During Reinitializing the Database

When you use `dbAdmin.bat` to reinitialize the database, warning messages resembling the following appear in the command prompt window that pops up when the `dbAdmin.bat` file is double-clicked. You can ignore these messages.

```
Connected.  
Creating FSRM DATA tablespace  
Creating FSRM INDX tablespace  
Connected.  
Warning: View created with compilation errors.  
Warning: View created with compilation errors.  
Warning: View created with compilation errors.
```

Database Password

The management server uses the following accounts to access and manage the database for the management server. You should change the passwords to these accounts to prevent unauthorized access.

Table 22 Default Passwords for Database Accounts

Account Name	Default Password
SYS (database admin user)	change_on_install
SYSTEM (database admin user)	manager
APPIQ_SYSTEM (management server admin user)	password

You should change the default password after the initial installation to prevent a security breach. See [“Changing the Password of System Accounts”](#) on page 151.

IMPORTANT: Do not change any Database Password manually or from the command line. If a password is changed manually, DB-Admin toll will not work and export, import, rman backup will start failing. See [“Changing the Password of System Accounts”](#) on page 151 for information about changing database passwords.

How to change the Database Password for APPIQ_USER account on Managed Database Application:

To discover and manage the database application from management server, APPIQ_USER (GLOBAL) account is used. On all the managed applications the APPIQ_USER account with the

default password of "password" was created when you configured the management server before you discovered the database application.

You can change the password for the APPIQ_USER account but this should be done for all database applications. See ["Changing the Password for the Managed Database Account"](#) on page 67.

For example: If the management server is managing two Oracle and two Sybase database applications, ask the database administrator to change the password for APPIQ_USER on all the Oracle and Sybase managed databases to a single password. Then, change the password for APPIQ_USER on the management server as described in ["Changing the Password for the Managed Database Account"](#) on page 67.

Troubleshooting Listener and Database Connection Problems

If there is any connection problem from management server or you see a JBoss connection problem in appiq.log, first verify that the listener and database are running:

Checking Listener Status:

If you are not able to start the listener contact technical support with the error message and network files (tnsnames.ora, listener.ora) from the following directory:

```
%ORA_HOME%\network\admin\listener.ora
```

- From the command line:

```
lsnrctl status
```

This will show you the status of listener. If the listener is not up try to start as:

```
lsnrctl start
```

- From the Database Admin Utility. See ["Checking the Database and Listener Status"](#) on page 151.

Checking Database Status:

These steps restart the database. If you receive an error message, contact technical support.

- From the command line:

```
sqlplus /nolog
```

```
sql>connect sys/change_on_install@appiq as sysdba
```

```
sql> startup force;
```

- From the Database Admin Utility. See ["Checking the Database and Listener Status"](#) on page 151.

These steps restart the database. If you receive an error message, contact technical support.


8 Viewing Element Topology and Properties

This chapter describes the following:

- ["About System Manager"](#) on page 161
- ["Accessing System Manager"](#) on page 163
- ["About Cisco Switches and VSANs in System Manager"](#) on page 163
- ["About the User Interface"](#) on page 164
- ["Viewing Storage Elements"](#) on page 182
- ["Setting Up Custom Commands"](#) on page 203
- ["Using External Tools"](#) on page 213
- ["About the Navigation Tab"](#) on page 214
- ["Viewing Element Properties"](#) on page 219
- ["Viewing Element Topology"](#) on page 222
- ["Creating a Virtual Application"](#) on page 231
- ["About the Provisioning Tab"](#) on page 232
- ["About the Events Tab"](#) on page 232
- ["Asset Attributes of an Element"](#) on page 233
- ["About the Collectors Tab"](#) on page 235
- ["About the Monitoring Tab"](#) on page 235
- ["About the Policies Tab"](#) on page 236
- ["Determining If a Host Belongs to a File System"](#) on page 236
- ["About the Data from CXFS File Systems"](#) on page 237

About System Manager

System Manager is the gateway to many features that let you view details about the discovered elements. System Manager provides a topology that lets you view how the devices in your network are connected. For example, direct attached storage connections are displayed by a dotted line.

NOTE: To view direct attach storage, you must enable the  button. See [Table 23](#) on page 164 for more information.

Use the utilities provided in the toolbar to modify the topology. For example, you can filter out fabrics and change the placement of elements in the topology through drag and drop functionality. See ["About the Toolbar in System Manager"](#) on page 164 for more information.

The following tabs provided in the middle pane provide additional information:

- **List** - Provides information about the elements by fabric and domain. See ["About the List Tab"](#) on page 167 for more information.

- **Access** - Provides information about zone entries, persistent bindings, and storage system LUN masking. You can also manage zone, zone aliases, and zone sets from this tab. See ["About the Access Tab"](#) on page 169 for more information.
- **Path** - Provides information about an element's path. See ["About the Path Tab"](#) on page 175 for more information.

When you right-click an element in the topology or in the List, Access, or Path tab, a drop-down menu appears. This menu provides additional functionality depending on the type of element clicked, such as telnet and the creation of zone sets. See the topics, ["About the Right-Click Menu Options"](#) on page 176.

If a switch has more than one connection to a host or storage system, the number of connections is displayed on the line connecting the elements. If there is only one connection, no number is displayed since the line indicates a connection exists.

Keep in mind the following:

- If your Java plug-in control panel cache setting is set to 50 MB, it is recommended you increase this setting to 100 MB or more. Increasing this setting improves the reloading performance of System Manager.
- Individual virtual SANs (VSANs) for Cisco switches are not displayed in the topology or fabric tree. The switches in a VSAN are displayed under the fabric to which their VSAN belongs. See ["About Cisco Switches and VSANs in System Manager"](#) on page 163 for more information.
- ISL connections are not shown as connected ports between two McDATA or Connectrix switches that are not both managed by EFC Manager or Connectrix Manager in a fabric.
- If multiple management servers are using the same Windows Proxy, hosts discovered by other management servers appear in the topology. The Windows Proxy tracks all Windows hosts discovered. If you have more than one management server using the same Windows Proxy, the management servers are aware of all hosts discovered by the Windows Proxy. For example, assume you have two management servers that are named A and B and are using the same Windows Proxy. You use management server A to discover host1, host2, and host3. You then use management server B to discover only host3. You will see host1, host2, and host3 in the topology for management server B, although host3 only appears in the lists for discovery, topology, and Discovery Data Collection.
- The following IBM HBA appears as QLogic HBAs in the Navigation and Properties pages, in addition to reports:
 - IBM MSJ
 - FaSt FC-2/2-133

By double-clicking an element in the topology you have access to the following features:

- **Navigation** - The Navigation tab provides information about an element and how it relates to other elements in its path. See ["About the Navigation Tab"](#) on page 214 for more information.
- **Properties** - The Properties tab provides a detailed status of the element. See ["Viewing Element Properties"](#) on page 219 for more information.
- **Topology** - The Topology tab provides a graphical representation of an element's path. It displays additional information not found in System Manager, such as adapters, slots, and fibre channel ports. See ["Viewing Element Topology"](#) on page 222 for more information.

- **Chargeback Manager** - The Asset Management tab lets you keep track of asset attributes, such as contact information for the element's owner. See "[Asset Attributes of an Element](#)" on page 233 for more information.
- **Collectors** - The **Collectors** tab lets you start a collector for a report and view the collector's corresponding reports once the information has been gathered. See "[About the Collectors Tab](#)" on page 235 for more information.
- **Provisioning Manager** - The Provisioning Manager tab lets you manage zones, zone sets, and zone aliases, in addition to pools, volumes, LUNs and LUN mappings. See the topics, "[About the Provisioning Tab](#)" on page 232 and "[About Provisioning Manager](#)" on page 239 for more information.
- **Events** - The Events tab lets you view events for an element. See the topics, "[About Event Monitoring for Storage Essentials](#)" on page 323 and "[About the Events Tab](#)" on page 232 for more information.
- **Policies** - The Policies tab lets you create utilization policies, which can send an e-mail, generate an event, or run a custom script when a set threshold for an element is triggered. See "[About Policy Manager](#)" on page 437 for more information.

Accessing System Manager

To access System Manager, click **Tools > Storage Essentials > System Manager** in HP Systems Insight Manager.

Keep in mind the following:

- If you are unable to access System Manager, make sure your Web browser is set to allow JavaScript and cookies.
- Java 2 Runtime Environment is required to access several features in the management server, such as System Manager. If you are accessing the management server and you do not have the Java 2 Runtime environment, you are asked to install it if your Web browser is on Windows. If your Web browser is on a Solaris system, you must manually install the Java plug-in. See "[Installing the Java Plug-in](#)" on page 8 for more information.
- When you are asked if you want to trust the signed applet for the software, click **Always**. The **Always** option prevents this message from being displayed every time you access System Manager.

About Cisco Switches and VSANs in System Manager

The management server does not display individual VSANs in its topology or fabric tree. The switches in a VSAN are displayed under the fabric to which their VSAN belongs. For example, assume switch_A belongs to VSAN1 and switch_B belongs to VSAN_2. VSAN_1 and VSAN_2 belong to the same fabric. The management server displays switch_A and switch_B under the same fabric without their VSAN listed in the tree.

To determine the VSAN to which a port on a Cisco switch belongs, access the Properties page for the port. The Properties page can be accessed by double-clicking the switch in System Manager and then clicking the **Properties** tab. Then, click the hyperlink for the port in the Properties page for the switch. You can also view information about Cisco port types, such as TE ports by accessing the Properties page for a port.

NOTE: If a TE port belongs to multiple VSANs, the management server mentions only the primary VSAN.

About the User Interface

This section describes the following:

- ["About the User Interface for System Manager"](#) on page 164
- ["About the Toolbar in System Manager"](#) on page 164
- ["Description of the Icons Displayed in the Topology"](#) on page 166
- ["About the List Tab"](#) on page 167
- ["About the Access Tab"](#) on page 169
- ["About the Path Tab"](#) on page 175
- ["About the Right-Click Menu Options"](#) on page 176

About the User Interface for System Manager

System Manager displays an easy to use interface, which provides the following:

- **Toolbar** - Provides buttons and menus to help you modify the topology in System Manager. See ["About the Toolbar in System Manager"](#) on page 164.
- **Tabs** - Provides information about individual elements. The following tabs are provided:
 - **List** - Provides information about the elements by type and by fabric and domain. See ["About the List Tab"](#) on page 167.
 - **Access** - Provides access to tools that let you provision and view information about provisioning. See ["About the Access Tab"](#) on page 169. Provisioning wizards may not be available in your kit. To determine if you can access Provisioning Manager, access the Feature List, which is accessible from the Documentation Center.
 - **Path** - Provides information about an element's path. See ["About the Path Tab"](#) on page 175.
- **Right-Click Menu** - Provides features you can use to manage that element. See ["About the Right-Click Menu Options"](#) on page 176.

About the Toolbar in System Manager

The following table provides a brief description of the buttons and menus in System Manager toolbar.

Table 23 Description of Toolbar in System Manager


Icon	Description
	Prints the topology. See "Printing Elements in System Manager" on page 187 for more information.

Table 23 Description of Toolbar in System Manager (continued)



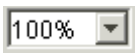







Icon	Description
	Magnifies the view
	Decreases the magnification
	Lets you set the magnification to a percentage of the default magnification
	Opens a smaller pane, which provides a global view of the topology. This lets you position the main view to a certain section of the topology. For more information, see "Using the Global View" on page 187.
	Fits the topology to the window, so you can see the entire topology.
	Lets you move an element in the topology. See "Arranging Elements in the Topology" on page 185.
	Lets you move multiple elements at once. This button is not accessible from the Topology tab. See "Arranging Elements in the Topology" on page 185.
	Lets you move the entire topology at once. Click the Pan button () and then click any place in the topology. Drag the mouse to a new location.
	<p>Updates the layout of the topology and removes the last saved layout from the database.</p> <p>Elements that have been manually moved might revert to their original position. This button is not accessible from the Topology tab.</p>

Table 23 Description of Toolbar in System Manager (continued)








Icon	Description
	<p>Saves the current topology, so that when you return to System Manager, the saved layout is restored.</p> <p>This option can be especially useful if you have moved elements in the topology and you want to keep their current location. This button is not accessible from the Topology tab.</p> <p>When you click the button, you are asked if you want the layout to apply to all users.</p> <ul style="list-style-type: none"> • Yes - All users who log into the management server see the topology you created. Only users with system configuration capability can save their layout for all other users • No - Other users cannot view the topology you saved. The saved topology appears the next time you log into the management server.
	<p>Opens a new window containing the topology. This feature lets you view different domains of the topology at once. This button is not accessible from the Topology tab.</p> <p>See "About the New Window Option" on page 228 for more information.</p>
	<p>Lets you view only selected fabrics in the topology. This button may not be accessible from the Topology tab. See "Filtering Fabrics" on page 194 for more information.</p>
	<p>Used in Capacity Manager to hide the lower pane.</p>
Search <input type="text"/>	<p>Lets you find an element by name or by Worldwide Name in the topology. You can type part of the information, and the management server highlights the elements that match.</p> <p>After you populate the search field, click the  button or press ENTER.</p>
Severity  <div> Warning Critical Major Minor Warning Informational Unknown </div>	<p>Displays the event severity icons for the elements displayed in the topology. See "Viewing Event Status in the Topology" on page 194 for more information. Disabled for Performance Manager and Capacity Manager.</p>








Table 23 Description of Toolbar in System Manager (continued)

Icon	Description
	This button calculates the topology path and host capacity. To view direct attached storage in System Manager, click this button. Direct attached storage is indicated by dotted lines.

Description of the Icons Displayed in the Topology

The following icons are displayed in the System Manager topology.

Table 24 Elements Displayed in the Topology

Icon	Description
	Indicates an application.
	Indicates a file server.
	Indicates a host. This particular icon is for a host running Microsoft Windows.
	Indicates a storage system.
	Indicates a switch. This particular icon is for a Brocade switch.
	Indicates a filer. This particular icon indicates a NAS filer.
	Indicates a tape library.

About the List Tab

The List tab provides information about the elements by type and by fabric and domain.

To find the fabrics in a domain, expand the domain node. You can see the elements in each fabric by expanding the fabric node, as shown in the following figure.




Figure 19 Expanding the Fabric Node

The “unknown” Fabric lists elements that have a Fibre Channel port connected to an undiscovered Fabric or the Fibre Channel port remains not connected.

When you click a fabric name in the tree, its members are highlighted in the right pane, as shown in the following figure.

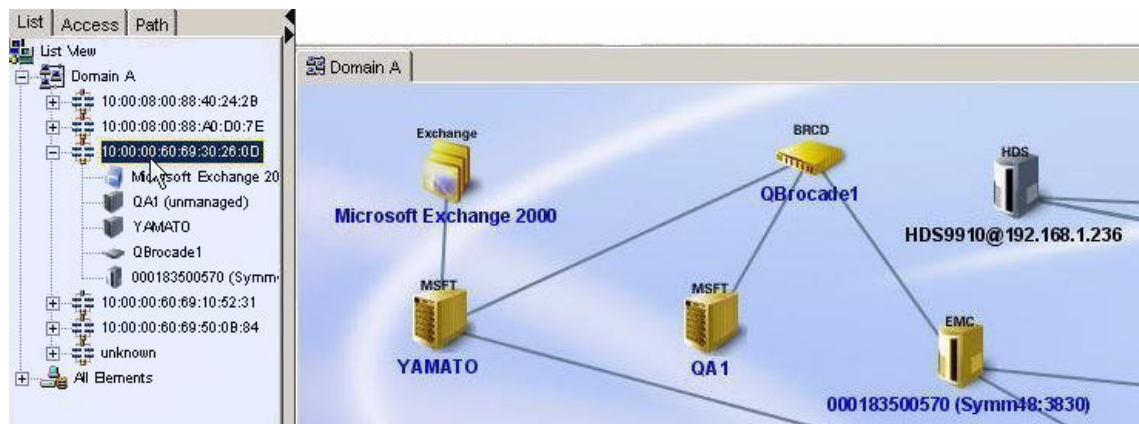


Figure 20 Fabric Members Are Highlighted

When you right-click an element in the List tab, a drop-down menu is displayed. The options displayed depend on the type of element you clicked. See “[About the Right-Click Menu Options](#)” on page 176 for an explanation of the options in the drop-down menu.

Viewing Elements by Type

You can view elements by type under the **All Elements** node on the **List** tab. This is especially helpful in determining how many elements you have of a specified type, such as the number of storage systems.

When you expand the element type node, all elements of that type are listed. If you select the element type node, all elements of that type are selected in the topology. For example, assume you want to determine the number of applications that the management server monitors. When you

expand the **Applications** node, the applications are listed. When you select the **Applications** node, the applications are highlighted in the topology, as shown in the following figure.

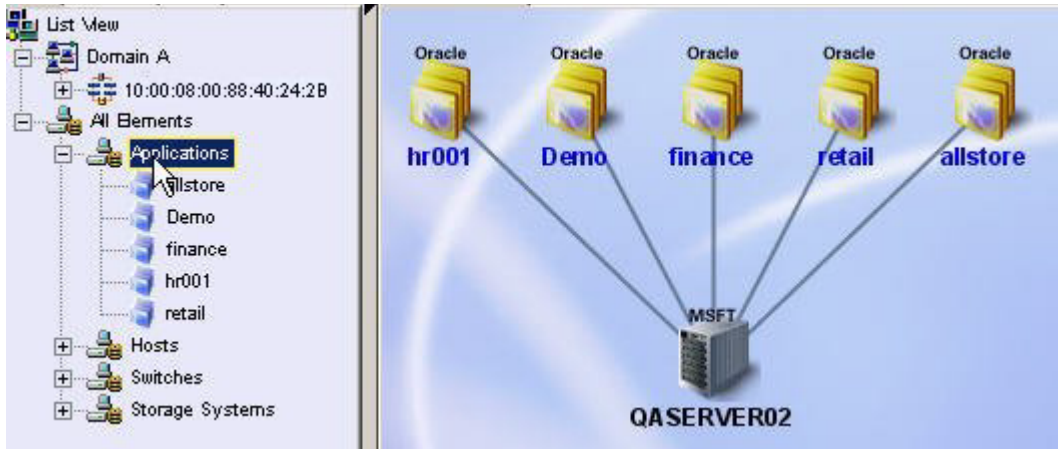


Figure 21 Highlighting the Applications in the Topology

If you select an element in the left pane, the element is highlighted in the topology. You also have access to additional functionality by right clicking the element. See "[About the Right-Click Menu Options](#)" on page 176 for more information.

About the Access Tab

The Access tab provides information about the following:

- Zone Entries
- Host Bindings
- Storage System LUN Masking
- You can also manage zone, zone aliases, and zone sets from this tab by right-clicking an element. See "[About the Right-Click Menu Options](#)" on page 176.

Obtaining Information About Zone Entries

To view the zone entries in a domain, expand **domain** > **Fabric name** > **Zone Set** nodes. Click the Zone Set node to see the members of the zone set highlighted in the right pane, as shown in the following figure.

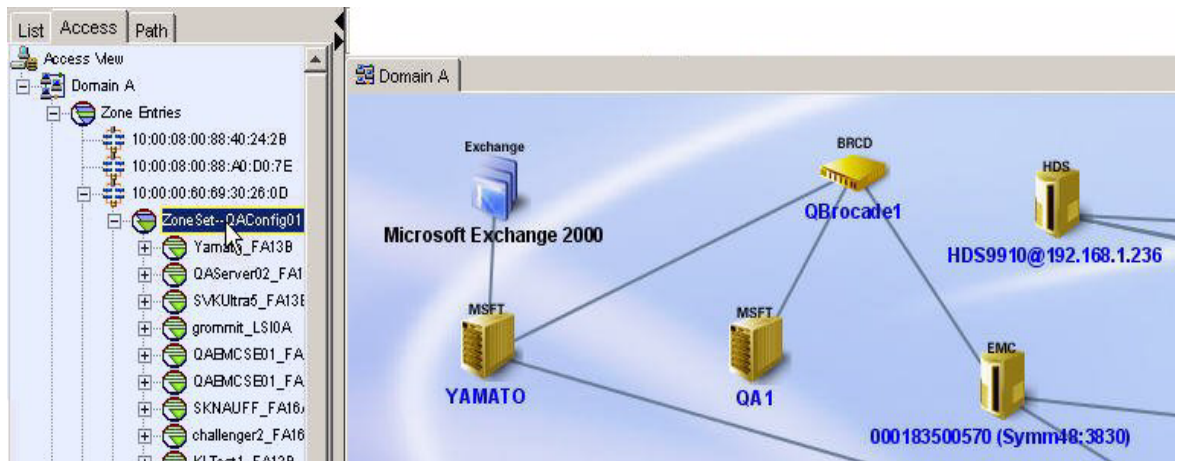




Figure 22 Members of a Zone Set

The  icon is displayed next to the name of the active zone set. The  icon is displayed next to the inactive zone sets.

To view members of a zone, do one or more of the following:

- Expand the node of the zone in the tree. The software displays the zone members underneath the node of the zone.

- Click the node of the zone in the tree. The software highlights the zone members in the right pane.

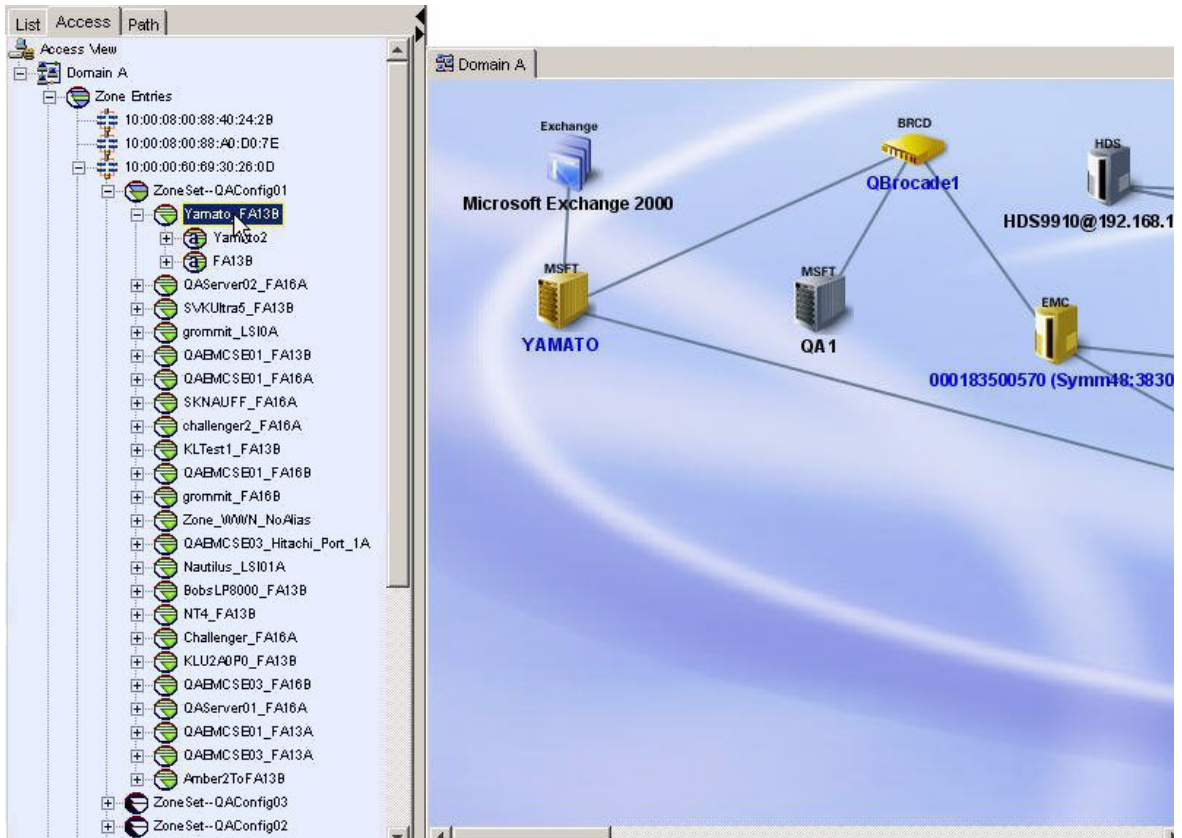


Figure 23 Displaying Zone Members

To view the relationship of the zone member to the switch, click the zone member in the tree. The software highlights the zone member and its switch.

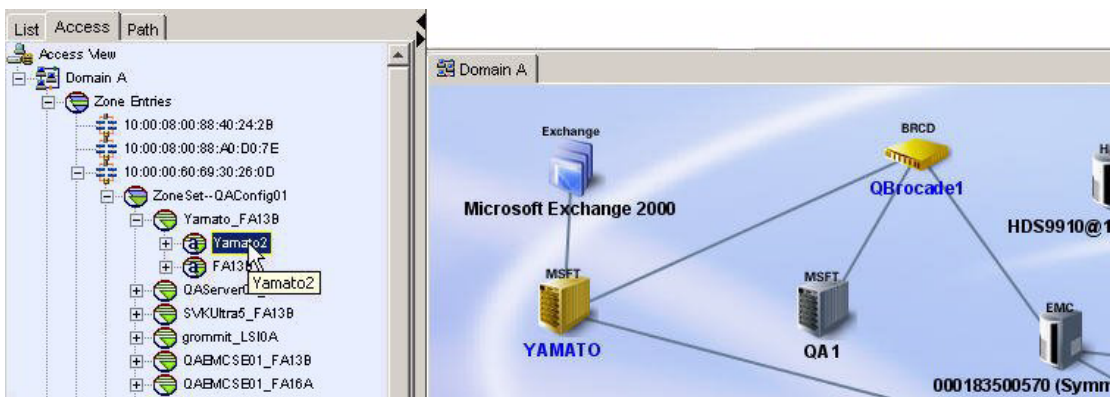


Figure 24 Zone Member to Switch

To view information about a zone member's port, expand the zone member node, as shown in the following figure. Notice that when you select the zone member node in the tree, it appears highlighted in the right pane.

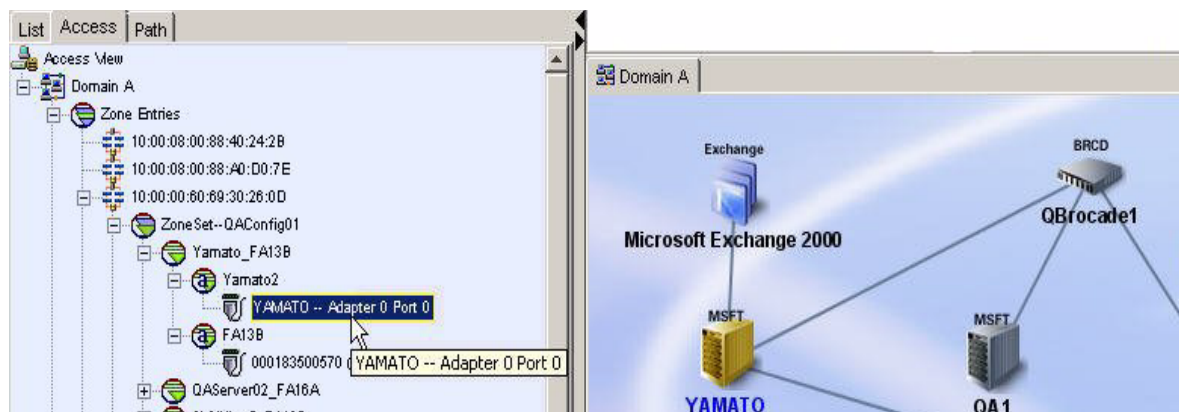


Figure 25 Obtaining Information About a Zone Member's Adapter

Obtaining Information About Host Bindings

To view the elements that have host bindings, click the Host Bindings node in the tree. The software highlights the elements that have host bindings in the right pane, as shown in the following figure.




Figure 26 Highlighting Elements with Host Bindings

To view a host binding, click the HBA node in the tree. The HBA node is under the element node, as shown in the following figure.

When you click the HBA node, the host and the element to which it has the binding are highlighted. A green line between the two elements indicates they have a binding.




Figure 27 Displaying Host Bindings

To view information about the ports on an HBA card, expand the HBA node in the tree, as shown in the following figure:



Figure 28 HBA Port Properties

Obtaining Information About Storage System LUN Masking

To obtain information about a storage system LUN masking, expand the fabric node. Then, expand the storage system node and click the Fibre Channel port. The values of the WWNs are displayed

under the node, and the storage system is highlighted in the right pane, as shown in the following figure.

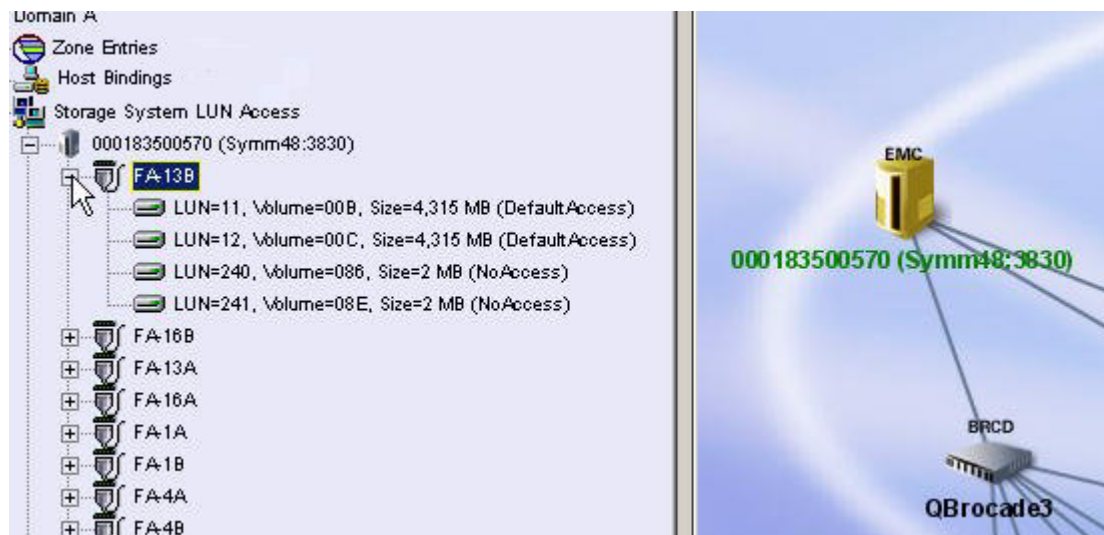


Figure 29 WWN Properties

The software displays properties of the WWN. If the LUN has a LUN masking, expand the LUN node to obtain information about the LUN masking.

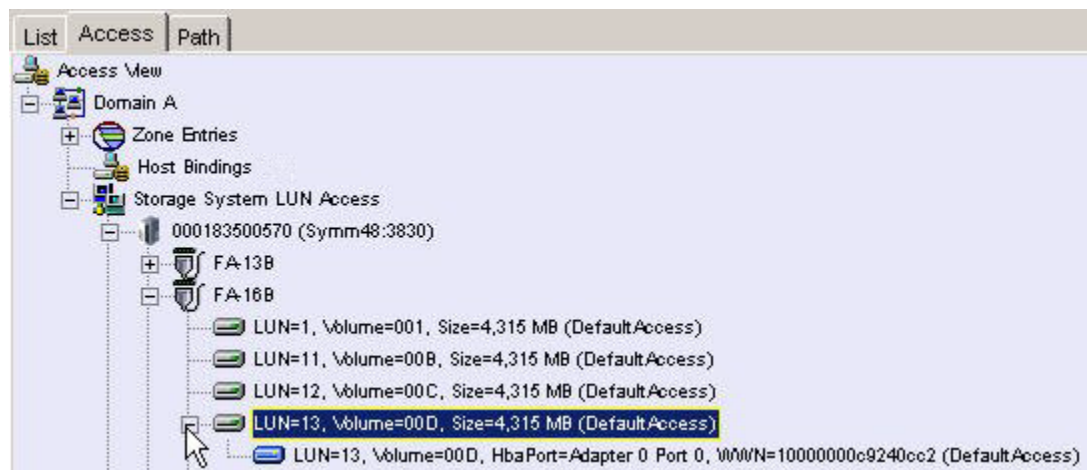


Figure 30 WWN Properties

To view a LUN masking, expand a LUN node.

About the Path Tab

The Path tab provides information about an element's path. By clicking a host's node, you can determine the host's path in the application. When you click a host node in the tree, the elements in the host's path appear highlighted in the right pane, as shown in the following figure.

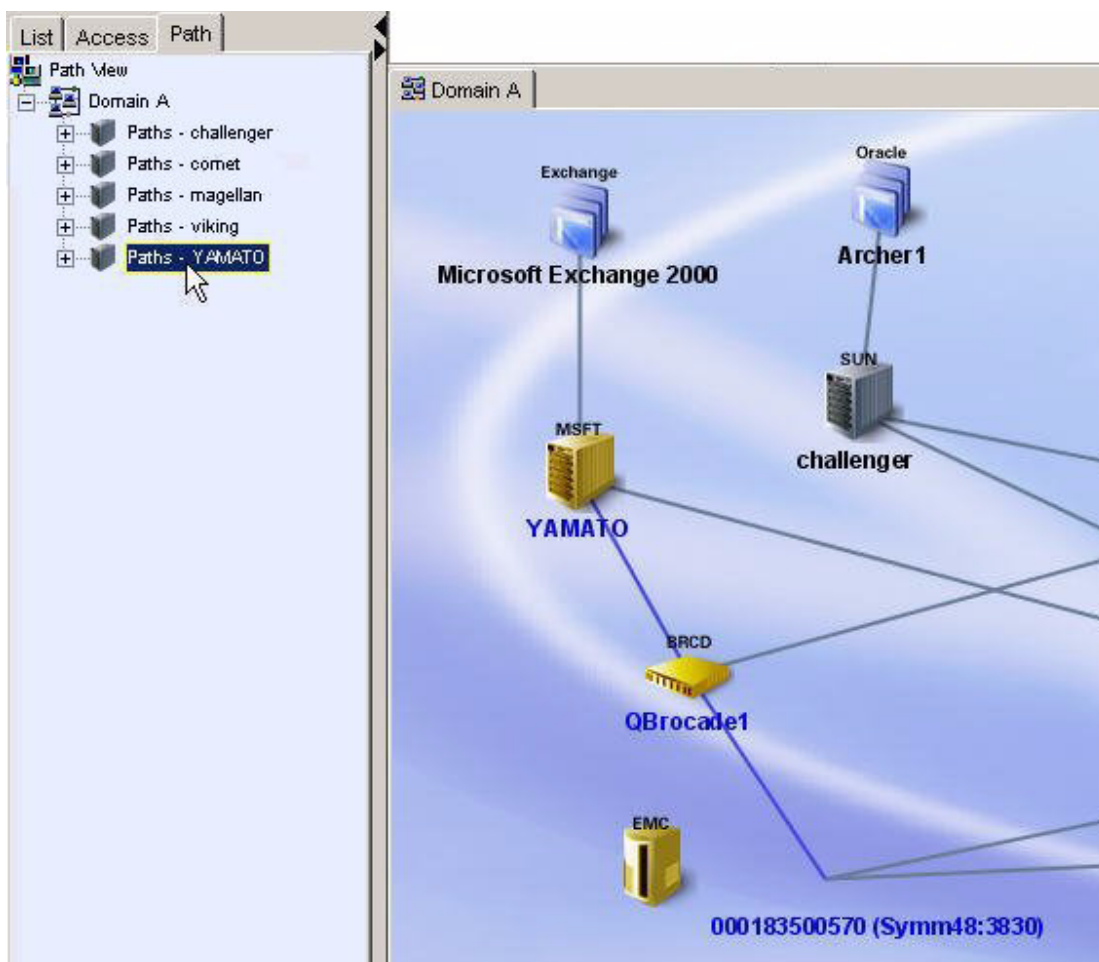


Figure 31 Obtaining Path Information

You can also determine the elements in a hosts path by expanding the **Application Path** and **Path** nodes under the host node in the tree.

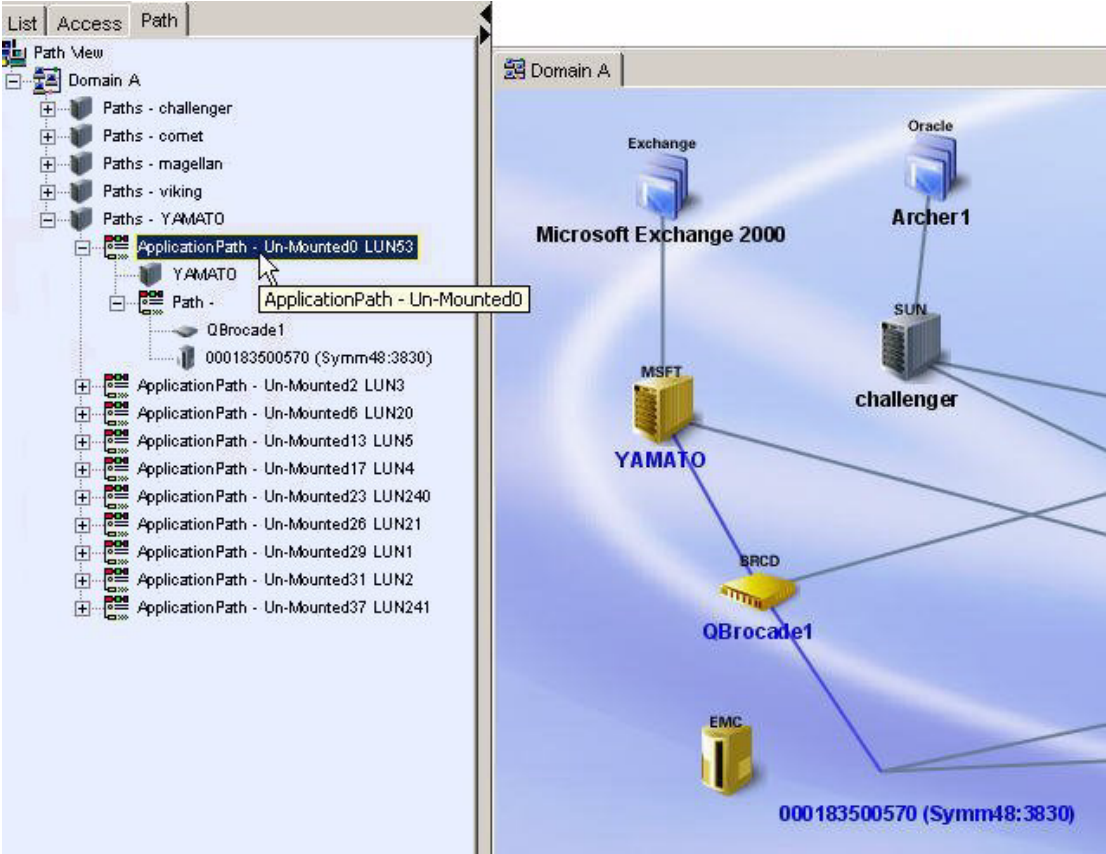


Figure 32 Path Information Visible in the Tree

When you right-click an element in the List tab, a drop-down menu is displayed. The options displayed depend on the type of element you clicked. See ["About the Right-Click Menu Options"](#) on page 176 for an explanation of the options in the drop-down menu.

About the Right-Click Menu Options

When you right-click an element in the topology pane or in the List or Access, or Path tabs, you see a drop-down menu. The options displayed in the menu depend on the type of element clicked.

See the following table for an explanation of the options displayed for elements right-clicked in the topology or right-clicked in the List or Path tab.

NOTE: Right-click menu options are not available to undiscovered fabrics.

Table 25 Drop-Down Menu Options Accessible from the Topology*

Menu Option	Description
Go to Navigation Details	This menu option redirects you to the Navigation page. If the element is labeled "Discovered," you are shown the Properties page. Elements are labeled "Discovered" when the management server has discovered the element, but it cannot obtain additional information about it. See " About the Navigation Tab " on page 214.
Go to Element Topology	Displays a graphical representation of the path of an element. This also includes multipathing. See " Viewing Element Topology " on page 222.
Show Events	Displays the events for an element. See " About the Events Tab " on page 232.
Show Policies	Displays the Policy tab for the element. You can then view, add, modify and delete policies assigned to the element. See " About the Policies Tab " on page 236 for more information.
Update Element Data	<p>The management server gathers new and changed details from the element and then redraws the topology with the updated information.</p> <p>Important:</p> <ul style="list-style-type: none">• Do not update element data during Discovery Data Collection. You can determine if the management server is getting the topology or all element details by looking at the label near the status button.• The Update Element Data functionality does not detect element components that have been removed, such as ports and LUNs. For example, assume you removed several LUNs from an array. If you right-click the storage system and select Update Element Data, the LUNs still appear in the user interface. You must perform Discovery Data Collection for the deleted LUNs to be removed from the user interface. See "Step 3 - Discovery Data Collection" on page 39.
Show Impact	Highlights the elements that are impacted. See " Showing the Impact of an Element " on page 190.
Show Port Details	Lets you determine the use of each port for all elements in the network. See " Viewing Ports " on page 190 for more information.

Table 25 Drop-Down Menu Options Accessible from the Topology* (continued)

Menu Option	Description
External Tools	<p>Provides several ways to access an element:</p> <ul style="list-style-type: none"> • Telnet - Lets you access a host or a switch through the telnet utility. The Telnet feature is only accessible to Web browsers on Microsoft Windows operating systems. • Browse - Lets you access the main Web page for a host or a switch. • Set Up External Tools - Lets you add URLs for accessing the management tools for the storage system. In some instances, the management tool for the storage system is accessible from this menu. For example, HiCommand for HDS storage systems and Command View for HP XP storage systems are accessible from the External Tools menu. <p>See "Using External Tools" on page 213 for more information.</p>
Discovered Elements	<p>Lets you group unnamed generic hosts. It provides the following options. See "About Hiding Generic Hosts" on page 201.</p> <ul style="list-style-type: none"> • Hide Generic Hosts for the Switch - Hides unnamed generic hosts that are connected to the switch. See "Hiding Generic Hosts for a Switch" on page 202. • Expand Generic Hosts for the Switch - Displays hidden unnamed generic hosts that are connected to the switch. See "Expanding Generic Hosts for a Switch" on page 202. • Hide Generic Hosts for All Switches - Hides unnamed generic hosts connected to all switches. See "Hiding Generic Hosts for All Switches" on page 203. • Expand Generic Hosts for All Switches - Displays hidden unnamed generic hosts that are connected to the switches. See "Expanding Generic Hosts for All Switches" on page 203.

Table 25 Drop-Down Menu Options Accessible from the Topology* (continued)

Menu Option	Description
Provision	<p>Provides provisioning tools for switches and storage systems:</p> <p>Switches - Lets you activate and deactivate zone sets, in addition to managing the following:</p> <ul style="list-style-type: none"> • Zone Aliases (Not applicable to McDATA switches) • Zones • Zone Sets (The delete zone set option is disabled for active zone sets). <p>Notes:</p> <ul style="list-style-type: none"> • When McDATA or Connectrix switches are discovered through a proxy by using SNMP, you cannot view or perform any provisioning operations for those switches. For example, you cannot view zone sets, zones, and/or zone alias. • When McDATA or Connectrix switches are discovered by their IP address by using SNMP, you can only view the active zone set and its members. You cannot create, modify, and/or delete zone sets or its members. <p>Storage Systems - Lets you manage the following:</p> <ul style="list-style-type: none"> • Storage Pools • Volumes • Host Security Groups <p>*Offered only when supported by storage system</p> <p>See "About the Provisioning Tab" on page 232, "SAN Zoning Overview" on page 240 and "Setting Up Storage Partitioning" on page 256 and for more information. Provisioning wizards may not be available in your kit. To determine if you can access Provisioning Manager, access the Feature List, which is accessible from the Documentation Center.</p>
Add Virtual Application	<p>Lets you add a virtual application so you can monitor it. A virtual application is a placeholder you create for an application. For example, you could create a virtual application for an application that was created just for your company. See "Adding a Virtual Application" on page 183.</p>
Set Business Cost (Applications only)	<p>Lets you assign a business cost to an application. See "Assigning a Business Cost to an Application" on page 192.</p>

Table 25 Drop-Down Menu Options Accessible from the Topology* (continued)

Menu Option	Description
Delete Element	Removes an element and its discovery instance from the system. It also removes other elements discovered through the removed element. See "Deleting Elements from the Management Server" on page 38. Important: If you are blocking pop-ups in a Netscape or Mozilla Web browser and you use the right-click menu to delete an element from System Manager, the Delete window is blocked and you are unable to delete the element. You must disable the popup blocker before you can delete the element.
Adding a Custom Command	Lets you run a custom command on an element, for example to start an executable or a script. See "Adding a Custom Command" on page 204.
Group together with other elements	Lets you group "Discovered" hosts and storage systems. See one of the following topics: <ul style="list-style-type: none"> • "Grouping "Discovered" Hosts" on page 197 • "Grouping "Discovered" Storage Systems" on page 198
Ungroup into multiple elements	Lets you ungroup "Discovered" hosts and storage systems. See one of the following topics: <ul style="list-style-type: none"> • "Ungrouping "Discovered" Hosts" on page 198 • "Ungrouping "Discovered" Storage Systems" on page 200
Recalculate Topology	Lets you know about topology changes. The management server contacts the elements on the topology list (Discovery > Topology) to determine topology changes. The management server uses this updated information to redraw the topology.
Change Fabric Name (Only Available from the List Tab)	Lets you change the name of the fabric. See "Changing the Fabric Name" on page 200 for more information.

*Additional menu items may appear for types of Automators and Advisors, such as Reachable Storage. The descriptions of the right-click menu options for a fabric appear in the following section, since these menu options are only accessible from the List tab.

When you right-click a fabric in the List tab, you are shown the following options:

- **Change Fabric Name** - Type the new fabric name in the Change Fabric Name window. Then, click **OK**.
- **Go to Properties** - Displays the properties of the fabric. See ["Viewing Element Properties"](#) on page 219 tab for more information. It also provides access to the Events and Provisioning Manager tabs. The Events tab displays events occurring within the fabric. The Provisioning

Manager tab lets you set up and manage zone provisioning. Provisioning Manager may not be available in your kit. To determine if you can access , access the Feature List, which is accessible from the Documentation Center.

- **Delete This Fabric** - Deletes a fabric. When you are asked if you want to delete the fabric, click **Yes** if you do not mind waiting for the management server to recalculate the topology. If the elements in the deleted fabric do not belong to another fabric, they are moved to the **unknown** node on the List tab.

Provisioning features are available from the right-click menu in the Access tab. When you right-click a zone, zone alias, or zone set in the Access tab, a drop-down menu is displayed. The provisioning options displayed in this menu depend on the type of element clicked.

IMPORTANT: If you do not see the provisioning features from the right-click menu, your license does not allow you to access provisioning.

The following table describes the menu options. To learn more about each task and its required steps, see "[Provisioning Manager](#)" on page 239.

Table 26 Drop-Down Menu Options on the Access Tab

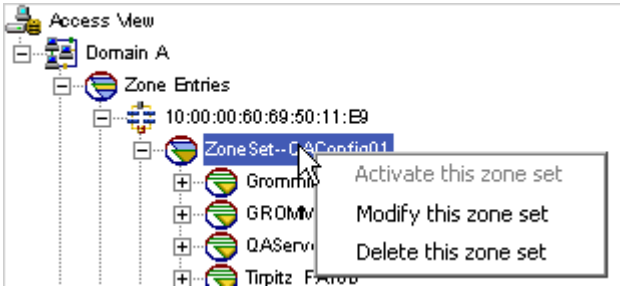



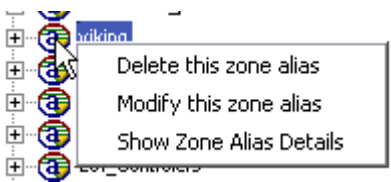
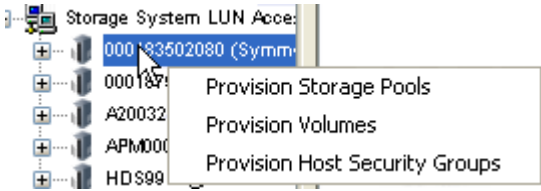
Task	To Perform Task, Right-Click...
1Activate a zone set	
1Modify a zone set	
1,3Delete a zone set	
1Create a zone	
1Delete a zone	
1Modify a zone	
1,2Show zone details	

Table 26 Drop-Down Menu Options on the Access Tab (continued)

Task	To Perform Task, Right-Click...
1Create a zone alias	A zone alias with a blue folder under the fabric node. 
1Delete a zone alias	A zone alias under the Zone Aliases node. 
1Modify a zone alias	
1Show zone alias details	
Provision Storage Pool	A storage system under the Storage System LUN Access node. 
Provision Volume	
Provision Host Security Groups	

1When McDATA and Connectrix switches are discovered through a proxy by SNMP, you cannot view or perform any provisioning operations for those switches.

2Only these options are accessible when McDATA and Connectrix switches are discovered by a switch IP address directly through SNMP. This setting provides view only access to the active zone set and its members. You cannot create, modify, and/or delete zone sets or its members.

3The delete zone set option is disabled for active zone sets.

Viewing Storage Elements

System Manager has a wide range of features to help you in viewing your storage elements in the topology. For example, you can filter fabrics, arrange elements in the topology, and search for elements in the topology. You can even view the status of elements in the topology and find the impact of removing an element.

See the following topics in this section for more information.

- [“Adding a Virtual Application”](#) on page 183
- [“Adding Information for “Discovered” Hosts”](#) on page 184
- [“Arranging Elements in the Topology”](#) on page 185
- [“Closing Topology Windows”](#) on page 187

- ["Using the Global View"](#) on page 187
- ["Printing Elements in System Manager"](#) on page 187
- ["Updating Element Data"](#) on page 189
- ["Viewing Ports"](#) on page 190
- ["Showing the Impact of an Element"](#) on page 190
- ["Showing the Impact of an Element"](#) on page 190
- ["Assigning a Business Cost to an Application"](#) on page 192
- ["Expanding the Topology Pane"](#) on page 194
- ["Filtering Fabrics"](#) on page 194
- ["Viewing Event Status in the Topology"](#) on page 194
- ["Custom Name for Switch Truncated in Topology"](#) on page 197
- ["Managing Groups"](#) on page 197
- ["Managing Fabrics"](#) on page 200
- ["Hiding and Showing Generic Hosts"](#) on page 201

Adding a Virtual Application

The management server lets you monitor applications not listed in the support matrix. For example, assume your company has created an internal application, and you want to be able to use the management server to monitor that application. You can create a virtual application for that product. A virtual application is a placeholder you create for an application.

NOTE: Only a user belonging to a role that has System Configuration selected on the **Edit Role** page (such as the Domain Administrator role) is allowed to create a virtual application.

Once you create the virtual application, it will appear connected to a host in your topology.

1. Type the following information for the virtual application.
 - **Name**
 - **Product**
 - **Description**
 - **Vendor**
 - **Version**
2. Select the storage volume for the application.

NOTE: You can view the properties of a volume by clicking its link.

3. Click **OK**.

The virtual application appears connected to the selected host.

Adding Information for “Discovered” Hosts

The software labels a host as “Discovered” when it cannot obtain additional information about a host it has discovered. To learn why the software was unable to obtain information about the element, see [“Troubleshooting Discovery and Discovery Data Collection”](#) on page 511.

If you have more than one “Discovered” host, it can be difficult to differentiate them. To make them easier to identify, you may want to add information about the host, such as the following:

- Custom Name
- IP Address
- DNS Name
- Operating System
- Version of the operating system


IMPORTANT: Do not add information for generic elements during Discovery Data Collection. You can determine if the management server is getting the topology or all element details by looking at label near the status button.

1. Click **Tools > Storage Essentials > System Manager** in HP Systems Insight Manager.
2. Double-click a “Discovered” host in the right pane.
The Properties tab is displayed.
3. In the **Custom Name** field, type a name for the element.
Keep in mind the following:
 - The name must contain 1 to 64 characters.
 - The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or one of the following symbols: dollar sign (\$), caret (^), hyphen (-), an underscore (_) or a space.
 - The name is case sensitive, for example, “Element1” and “element1” are different elements.
4. In the **IP Address** field, type an IP address for the element.
5. In the **DNA Name** field, type a DNS name for the element.
6. In the **Version** field, type the version of the operating system.
7. In the **Operating System** field, type one of the following operating systems:
 - **AIX** - corresponds to IBM AIX®
 - **HP-UX** - corresponds to all versions of HP-UX™
 - **IRIX** - corresponds to SGI IRIX®
 - **Linux**
 - **Windows** - corresponds to Microsoft Windows®
 - **Solaris** - corresponds to Sun Solaris™
8. Click **Save**.


When you access System Manager, the information you entered appears in the topology.

Arranging Elements in the Topology


To improve usability, arrange the topology so it suits your environment. For example, if you plan to filter various fabrics, you might want to arrange the topology so that elements are arranged by fabric. This way when you filter the fabrics, large gaps do not appear in the topology. You can arrange elements individually or in groups, as described in the following sections.

NOTE: The topology displays direct attached connections as a dotted line from the host to the storage system. To view direct attach storage, you must enable the  button. See [Table 23](#) on page 164 for more information.


To arrange elements individually:

1. Click the  button.
2. Click the element you want to move and drag it to a new location.
3. Repeat the previous step for each element you want to move.

The management server provides buttons to help you with viewing and arranging the topology. To learn more about those buttons, see "[About the Toolbar in System Manager](#)" on page 164.

4. Once you have finished arranging the topology, click the  button to save it.
5. To learn more about filtering fabrics, see "[Filtering Fabrics](#)" on page 194.

To arrange elements in a group:

1. Click the  button.
2. Holding down the mouse button, drag the mouse diagonally across the elements you want to move. Then, lift your finger off the mouse button when you are done.

A square encloses the elements, as shown in the following figure. If you want to redo the square, just click outside of the square and retry.



Figure 33 Enclosing the Elements

3. To move the elements within the square, click within the square. Holding down the mouse button, drag the elements to the new location.




Figure 34 Dragging Multiple Elements to Their New Location


Closing Topology Windows


Whenever you select a new topology view, the software creates a pane for that view.

To lessen the number of panes open, do the following:

1. Right-click the tab of one of the views.
2. Select one of the following from the drop-down menu:
 - **Close** - Closes the current topology pane in System Manager.
 - **Close All** - Closes all of the topology panes in System Manager.
 - **Close All But Current** - Closes all of the topology panes in System Manager, except the current one.

Using the Global View



If you have a large storage network, navigating it can be daunting, especially if it cannot fit in the pane. The global view () provides “a bird's eye” view of the network. With this view, you can move the viewing area to a certain section of the network.

1. Click the  button at the top of the screen.
A smaller pane displaying “a bird's eye” view of topology appears.
2. Move the brackets so they enclose the area of the network you want to view in the main pane.

Printing Elements in System Manager

You can print the topology displayed in System Manager.

To print the topology:

1. Click **Tools** > **Storage Essentials** > **System Manager** in HP Systems Insight Manager.
2. If the topology spans more than one screen, arrange the elements so they are closer together, preferably on one screen. To move an element, click the  button and then click the element you want to move. Drag the element to its new location. Moving elements closer together prevents the printout from appearing too stretched.
3. Click the  button.
4. Use the fields on the Paper tab to modify the setup of the page. When you are done, click **Apply**. If you want the default settings, click **Default**.
A preview of the printout is displayed in the right pane.

IMPORTANT: Before you change the margins, decide on a unit of measurement.

- **Paper format** - Select the paper size from the drop-down menu.
- **Unit** - Select cm (centimeters) or inch for the margins.
- **Paper width** - Displays the width of the paper. You can modify the measurement in this field when you select the **Custom** option in the Paper format drop-down menu.

- **Paper height** - Displays the height of the paper. You can modify the measurement in this field when you select the **Custom** option in the Paper format drop-down menu.
 - **Top margin** - Type a measurement for the top margin.
 - **Bottom margin** - Type a measurement for the bottom margin.
 - **Left margin** - Type a measurement for the left margin.
 - **Right margin** - Type a measurement for the right margin.
 - **Orientation** - Click an orientation for the printout.
5. Click the **View Selection** tab to modify how the printout will appear on the page. You can modify the following. When you are done, click **Apply**. If you want the default settings, click **Default**.
- A preview of the printout is displayed in the right pane.

IMPORTANT: Before you change the margins, decide on a unit of measurement.

- **Start x** - Determines the horizontal placement of the printout on the page with zero being the closest to the right margin. For example, if the value is 50 for **Start x**, the printing starts at 50 inches or centimeters (depending on what you selected) from the right margin. You can also enter negative numbers. Anything more than zero expands the printout to another page.
 - **Start y** - Determines the vertical placement of the printout on the page with zero being the closest to the bottom margin. For example, if the value is 50 for **Start y**, the printing starts at 50 inches or centimeters (depending on what you selected) from the bottom. You can also enter negative numbers.
 - **Width** - Determines the width of the printout.
 - **Height** - Determines the height of the printout.
- To remove extra space around the topology, click the **Trimmed** button.
6. To change how many pages the printout will use, select one of the following. When you are done, click **Apply**. If you want the default settings, click **Default**.
- A preview of the printout is displayed in the right pane.

IMPORTANT: Before you change the margins, decide on a unit of measurement.

- **Unit** - Select cm (centimeters) or inch for the margins.
- **Position/Size** - Lets you change the position and size of the printout so that it spans several pages:
 - **Start x** - Determines the horizontal placement of the printout on the page with zero being the closest to the right margin. For example, if the value is 50 for **Start x**, the printing starts at 50 inches or centimeters (depending on what you selected) from the right margin. Determines the horizontal placement of the printout. Anything more than zero expands the printout to another page.

- **Start y** - Determines the vertical placement of the printout on the page with zero being the closest to the bottom margin. For example, if the value is 50 for **Start y**, the printing starts at 50 inches or centimeters (depending on what you selected) from the bottom.
 - **Width** - Determines the width of the printout. If the width entered does not fit on the page, the printout wraps around to another page.
 - **Height** - Determines the height of the printout. If the height entered does not fit on the page, the printout wraps around to another page.
 - **Resolution (pixel/unit)** - Lets you change the resolution so that the printout spans several pages.
 - **Page** - Lets you expand the printout so it prints on several pages without modifying the graphic.
7. To preview your pages, click the Preview tab. Then click the page you want to preview. The page appears in the right pane.
 8. When you are ready to print, click **Print**.
 9. Click **Close**.

NOTE: To revert back to all of the original settings, click the **Default** button next to the **Print** button.

Updating Element Data

System Manager lets you update data about elements directly from this screen. When you update element data, the management server updates infrastructure details from the element and then redraws the topology with the updated information.

Keep in mind the following:

- Do not update element data during Discovery Data Collection. You can determine if the management server is getting the topology or all element details by looking at the label near the status button.
- Updating element data does not detect element components that have been removed, such as ports and LUNs. For example, assume you removed several LUNs from an array. If you right-click the storage system and select **Update Element Data**, the LUNs still appear in the user interface. You must perform Discovery Data Collection for the deleted LUNs to be removed from the user interface. See "[Step 3 - Discovery Data Collection](#)" on page 39.
- Update Element Data does not correctly update switch ISL and topology changes. Discovery Data Collection must be performed to obtain switch ISL and topology changes.

To update an element:

1. Click **Tools** > **Storage Essentials** > **System Manager** in HP Systems Insight Manager.
2. Right-click the element you want to update.
3. Select **Update Element Data** from the drop-down menu.
The software begins to update its database with the updated infrastructure details from the element.

When the process is complete, the status button returns to green.

Viewing Ports

When looking at an element on the network, such as a switch, it can be difficult to determine how the ports are used. System Manager provides a view that lets you determine the use of each port for all elements in the network.

To view the ports:

1. Do one of the following:
 - **Access System Manager** - Click **Tools > Storage Essentials > System Manager** in HP Systems Insight Manager.
 - **Access the Topology page (for an element)** - Do one of the following:
 - Double-click an element in System Manager, and then click the **Topology** tab
 - Right-click an element, and then select **Show Element Topology** from the drop-down menu.
2. Right-click an element in the topology.
3. From the drop-down menu, select **Show Port Details**.
The ports are displayed.

Showing the Impact of an Element

You can display an element's impact. For example, assume you want to replace a switch. You can use this feature to determine which elements in the network would be impacted by taking the switch off the network.

To find an element's impact:

1. Right-click the element from which you want to obtain impact information.
2. Select **Show Impact** from the drop-down menu.

The other elements in the path of the element you right clicked are highlighted in yellow. For example, assume you right-clicked the Oracle instance klu2e (shown in the following figure) and selected **Show Impact** from the drop-down menu. The elements on which klu2e is

dependent are highlighted in yellow. This means that if any of these highlighted elements are removed from the network, klu2e may have difficulty functioning.




Figure 35 Showing the Impact of an Application

However, the show impact feature not only shows the other elements on which an element is dependent, but it also displays the other elements dependent on it. For example, assume you right-clicked a switch and selected **Show Impact** from the drop-down menu. The highlighted elements would include its dependent elements, such as the hosts, applications, and storage systems connected to it. These elements might difficulty communicated with one another if the switch was removed.

Likewise, if you decided to show the impact of a host, the highlighted elements would not only include its dependent elements, such as its applications, but also the elements on which it is dependent, such as switches.

Use the following table as a guideline to help you in determining whether the highlighted elements are dependent or required.

Table 27 Show Impact Results

If you select "Show Impact" for	Highlighted Elements
An Application (virtual or real)	Elements required by the application, such as its host and a switch.
A Host	<ul style="list-style-type: none">• Elements dependent on the host, such as its applications.• Elements required by the host, such as switches.
A Switch	Elements dependent on the switch, such as hosts and storage systems
Storage System	Elements dependent on the storage system, such as hosts.

Assigning a Business Cost to an Application

The management server lets you assign a business cost to an application, including virtual applications. This information is used in Event Monitoring for Storage Essentials for ranking events from elements. Event Monitoring for Storage Essentials determines the rank of an event by taking into account the business cost of the application and the severity of the event. You can sort events by rank in Event Monitoring for Storage Essentials by clicking the Rank column.

For example, assume you assigned a business cost of \$40 to one application and a business cost of \$30 to another application. If an event occurs with the same severity level occurs from both applications. The \$40 application has a higher rank because it has a higher "business cost."

The events from the elements in the path of the applications also inherit a business cost from the applications that use it. For example, assume a host has an application assigned a cost of \$30. The host would have a business cost of \$30. If the host has two applications (both valued at \$30), the host would be valued at \$60 because the two applications are using it. Likewise the switch connected to the host would also have a value of \$60 because the two \$30 applications use it. If a switch has a \$40 application on one host and two \$30 applications on another host, the switch has a value of \$100, see the following figure.

The cost of the storage system is determined by the applications in its path. Two storage systems connected to a switch can have different business costs based on the applications in their path. For

example, a storage system would have a value of \$60 if two \$30 applications are in its path, as shown in the following figure.




Figure 36 Determining Business Cost

To assign a business cost to an application:

1. Do one of the following:
 - Right-click an application in System Manager, and then select **Set Business Cost** from the drop-down menu.
 - Double-click an application in System Manager. Click the **Properties** tab. Then click the **Change** button next to the Business Cost field.
 - Click an application in Application Viewer. Click the **Properties** tab. Then click the **Change** button next to the Business Cost field.
2. In the **Business Cost** field, type an amount, for example 35.25.
3. If the host has access to storage volumes, click the **Volumes** tab. Then, select the volumes the application uses.
4. Click **OK**.

Expanding the Topology Pane

To increase screen space for viewing the topology, hide the List, Access, and Path tabs by clicking the arrow pointing left on the border between the pane containing the tabs and the main pane in System Manager.




Figure 37 Expanding the Topology Pane

To display these tabs, click the arrow pointing right on the border for the left pane.

Filtering Fabrics

To view a specified fabric in the topology, do the following:



1. Click **Tools > Storage Essentials > System Manager** in HP Systems Insight Manager.
2. Click the  button near the top of the screen.
3. Deselect the fabrics you do not want to view, as shown in the following figure:



Figure 38 Filtering Fabrics

4. Click the **Apply Filter** button.
System Manager displays the selected fabrics.

Viewing Event Status in the Topology

You can obtain a status of the events occurring on the elements displayed in System Manager by clicking the  button located on the toolbar. Elements with events that have occurred within the

last five minutes have a triangle displayed next to them. The color of the triangle corresponds to the severity of the event.









NOTE: The Event Status button () is disabled in Capacity Manager and Performance Manager.

Table 28 Legend for Severity Level Icons

Icon	Description
	The event has a critical impact.
	The event has a major impact.
	The event has a minor impact.
	The event is providing a warning.
	The event is providing information.
	The severity of the event is not known.

The triangle corresponding to the highest severity is shown. For example, the management server displays an icon for a critical event next to an element that might also have minor events.

Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

Use the Severity drop-down menu to filter which type of events you want to view. It displays the severity icons with the selected severity level or higher. For example, you can be notified of just critical and major events by selecting **Major** from the Severity drop-down menu and clicking the  button.

See the following table for a description of each severity level.

Table 29 Descriptions of the Severity Levels Presented in the Severity Menu

Severity Level	Description
Critical	Denotes elements that have a critical severity level. For example, Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

Table 29 Descriptions of the Severity Levels Presented in the Severity Menu (continued)

Severity Level	Description
Major	<p>Denotes elements that have a major severity level. It also displays icons for the critical severity level.</p> <p>For example, one or more physical fabric objects (device port, switch, or fabric) have disappeared.</p>
Minor	<p>Denotes elements that have a minor severity level. It displays icons for the following severity levels:</p> <ul style="list-style-type: none"> • Critical • Major <p>An example of a minor event is a physical fabric object (switch port or fabric) has changed state.</p>
Warning	<p>Denotes elements that have a warning severity level. It displays icons for the following severity levels:</p> <ul style="list-style-type: none"> • Critical • Major • Minor <p>An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.</p>
Informational	<p>Denotes elements that have an informational severity level. It displays icons for the following severity levels:</p> <ul style="list-style-type: none"> • Critical • Major • Minor • Warning <p>An example of an informational event is a progress report event for firmware download operation currently in progress.</p>
Unknown	<p>Denotes elements that have unknown severity level. It displays icons for the following severity levels:</p> <ul style="list-style-type: none"> • Critical • Major • Minor • Information

Custom Name for Switch Truncated in Topology

If a custom name for a switch is long, its name may appear truncated in the topology. The full name appears once the cursor is positioned over the switch with the custom name.

Managing Groups

This section describes the following:

- ["About Groups"](#) on page 197
- ["Grouping "Discovered" Hosts"](#) on page 197
- ["Ungrouping "Discovered" Hosts"](#) on page 198
- ["Grouping "Discovered" Storage Systems"](#) on page 198
- ["Ungrouping "Discovered" Storage Systems"](#) on page 200

About Groups

System Manager lets you group together hosts and storage systems that have been labeled "Discovered" so the management server sees them as one element in the topology. The management server labels an element as "Discovered" when it has discovered the element, but it cannot obtain additional information about it. Grouped elements preserve space in the topology since only one element is displayed to represent the group. It also provides a way to keep track of all your "Discovered" hosts and storage systems.

Grouping "Discovered" Hosts

If you have several "Discovered" hosts, you might want to group them together so the management server sees them as one element in the topology. The management server labels a host as "Discovered" it when it has found the host but it cannot obtain additional information about it. Grouped elements preserve space in the topology since only one "Discovered" host is displayed to represent the group. Grouping also provides a way to keep track of your "Discovered" hosts.

Keep in mind the following:

- A user's role must include an access level of Element Control or Full Control for hosts. See the topic, ["Editing Roles"](#) on page 84 for more information about the access level of a role.
- Grouped elements are still seen as "Discovered" so the management server is unable to monitor or manage them.
- Do not create groups during "Discovery Data Collection." You can determine if the management server is getting the topology or all element details because the status button appears red during both operations.
- You can tell if a host is generic by double-clicking the host in System Manager and then clicking the Properties tab. If a host is generic, it is listed as "Generic Host" for its description.

To group "Discovered" hosts:

1. Access System Manager, as described in ["Accessing System Manager"](#) on page 163.
2. Right-click a "Discovered" host and select **Group together with other elements** from the drop-down menu.
3. In the **Custom Name** field, type a custom name for the group.

4. In the **IP Address** field, type an IP address for the group.
5. In the **DNS Name** field, type the DNS name for the group.
6. In the **Version** field, type a version number for the group.
7. In the **Operating System** field, type the operating system for the hosts in the group.
8. Select the hosts you want to be a part of the group and click the button with the greater than sign (>).

The hosts are added to the group.

You can sort the hosts by:

- **Name** - To sort hosts by name, click the Hosts column heading.
- **Port** - To sort hosts by port, click the Ports column heading.
- **Connected Switches** - To sort hosts by connected switches, click the Connected Switches column heading.

An arrow appears next to the column heading that sorts the hosts. For example, if the hosts are being sorted by name, an arrow appears next to the Hosts column heading. If the arrow next to the column heading is pointing up, the hosts are sorted in ascending numerical and alphabetical order. If the arrow is pointing down, the hosts are sorted in descending numerical and alphabetical order. Click the column heading consecutively to change the direction of the arrow.

9. To remove hosts from the group, click the button with the less than sign (<).

10. Click **OK**.

The management server no longer displays the grouped elements in the topology individually. A host icon with the group name on the bottom represents the group. The group cannot be monitored or managed.

Ungrouping "Discovered" Hosts

You might want to ungroup "Discovered" hosts, if one of the hosts in the group is going to change. For example, the host will be taken off line.

Keep in mind the following:

- Do not ungroup elements during "Discovery Data Collection." You can determine if the management server is getting the topology or all element details because the status button appears red during both operations.
- A user's role must include an access level of Element Control or Full Control for hosts. See the topic, "[Editing Roles](#)" on page 84 for more information about the access level of a role.

To ungroup multiple elements:

1. Access System Manager, as described in "[Accessing System Manager](#)" on page 163.
2. Right-click host icon for a group and select **Ungroup into multiple elements** from the drop-down menu.
3. When you are asked if you want to ungroup the elements, click **OK**.

The elements are ungrouped.

Grouping "Discovered" Storage Systems

If you have several "Discovered" storage systems, you might want to group them together so the management server sees them as one element in the topology. The management server labels a storage system as "Discovered" when it has found the storage system, but it cannot obtain additional information about it. Grouped elements preserve space in the topology since only one "Discovered" element is displayed to represent the group. It also provides a way to keep track of all your "Discovered" storage systems.

Keep the following in mind:

- A user's role must include an access level of Element Control or Full Control for storage systems. See ["Editing Roles"](#) on page 84 for more information about the access level of a role.
- Grouped elements are still seen as "Discovered", so the management server is unable to monitor or manage them.
- Do not group storage systems during "Discovery Data Collection." You can determine if the management server is getting the topology or all element details by looking at label near the status button.
- You can tell if a storage system is generic by double-clicking the host in System Manager and then clicking the Properties tab. Look for its description.

To group "Discovered" storage systems:

1. Access System Manager, as described in ["Accessing System Manager"](#) on page 163.
2. Right-click a "Discovered" storage system and select **Group together with other elements** from the drop-down menu.
3. In the **Custom Name** field, type a custom name for the group.
4. In the **Vendor** field, type the vendor names for the storage systems in the group.
5. Select **Tape Library** if you want the entire group to be considered a tape library.
Keep in mind this tape library will be considered as discovered, meaning it will not be managed or monitored by the management server.
6. Select the storage systems you want to be a part of the group and click the button with the greater than sign (>).

The storage systems are added to the group.

You can sort the storage systems by:

- **Name** - To sort storage systems by name, click the Storage Systems column heading.
- **Port** - To sort storage systems by port, click the Ports column heading.
- **Connected Switches** - To sort storage systems by connected switches, click the Connected Switches column heading.

An arrow appears next to the column heading that sorts the storage systems. For example, if the storage systems are being sorted by name, an arrow appears next to the Storage Systems column heading. If the arrow next to the column heading is pointing up, the storage systems are sorted in ascending numerical and alphabetical order. If the arrow is pointing down, the storage systems are sorted in descending numerical and alphabetical order. Click the column heading consecutively to change the direction of the arrow.

7. To remove storage systems from the group, click the button with the less than sign (<).

8. Click **OK**.

The management server no longer displays the grouped elements in the topology individually. A storage system icon with the group name on the bottom represents the group. The group cannot be monitored or managed.

Ungrouping "Discovered" Storage Systems

You might want to ungroup "Discovered" storage systems, if one of the storage systems in the group is going to change. For example, the storage system will be taken off line.

Keep in mind the following:

- A user's role must include an access level of Element Control or Full Control for storage systems. See the topic, "[Editing Roles](#)" on page 84 for more information about the access level of a role.
- Do not ungroup elements during "Discovery Data Collection." You can determine if the management server is getting the topology or all element details by looking at label near the status button.

To ungroup multiple elements:

1. Access System Manager, as described in "[Accessing System Manager](#)" on page 163.
2. Right-click a storage system icon for a group and select **Ungroup into multiple elements** from the drop-down menu.
3. When you are asked if you want to ungroup the elements, click **OK**.
The elements are ungrouped.

Managing Fabrics

This section describes the following:

- "[Changing the Fabric Name](#)" on page 200
- "[Deleting Fabrics](#)" on page 200

Changing the Fabric Name

To change a fabric name:

1. Access System Manager by clicking the **System Manager** button in the left pane.
2. Click the **List** tab.
3. Right-click a fabric name.
4. Select **Change Fabric Name** from the drop-down menu.
5. In the **Enter a Fabric Name** field, type a new fabric name.
6. Click **OK**.

Deleting Fabrics

When you delete a fabric, the elements in the fabric are not removed. After you delete the fabric, the management server recalculates the entire topology. The recalculation may take some time, especially if you have a large topology.

To delete a fabric:

1. Access System Manager by clicking the **System Manager** button in the left pane.
2. Click the **List** tab.
3. Right-click a fabric name.
4. Select the **Delete This Fabric** option from the drop-down menu.
5. When you are asked if you want to delete the fabric, click **Yes**.

The management server recalculates the topology. If the elements in the deleted fabric do not belong to another fabric, they are moved to the **unknown** node on the List tab.

Hiding and Showing Generic Hosts

This section describes the following:

- ["About Hiding Generic Hosts"](#) on page 201
- ["Hiding Generic Hosts for a Switch"](#) on page 202
- ["Expanding Generic Hosts for a Switch"](#) on page 202
- ["Hiding Generic Hosts for All Switches"](#) on page 203
- ["Expanding Generic Hosts for All Switches"](#) on page 203

About Hiding Generic Hosts

You can reduce the amount of time it takes to arrange your topology, by using the hide unnamed generic hosts feature to hide hosts that a switch has detected. An element is considered to be generic if the management server can detect the element but it cannot obtain additional information about the element during "Getting the Topology" or "Discovery Data Collection."

When you use the show/hide feature, your changes persist to the next time you log into the management server. If you log in as another user, you will not see your changes. This feature allows each user to arrange the topology as he or she wishes.

NOTE: This feature only pertains to unnamed generic hosts. If you name a generic host, you cannot use this feature to hide the named host. The hide generic element feature also does not work for grouped unnamed generic hosts and missing elements. To learn how to give a custom name to an unnamed generic host, see ["Assigning a Custom Name"](#) on page 221. To learn more about groups, see the topic, ["About Groups"](#) on page 197.

The management server provides two variations of this feature:

- **Hiding Generic Hosts:** This feature hides unnamed generic hosts detected by a switch. The management server detects an element by looking at the ports on a switch. If it cannot find additional information about the element, it marks it as generic with a question mark over the element.

For example, assume you have a switch with 10 "discovered" elements. The management server detected these elements by looking at the ports on the switch and determined the type of element connected. "Discovered" elements appear with a question mark above their icon in the topology. The question mark indicates the management server has detected the element, but it

cannot obtain additional information. To learn how to use this feature, see ["Hiding Generic Hosts for a Switch"](#) on page 202.

- **Hide Generic Hosts for All Switches:** This feature hides unnamed generic hosts within a domain. To learn how to use this feature, see ["Hiding Generic Hosts for All Switches"](#) on page 203.

Hiding Generic Hosts for a Switch

Simplify your topology by hiding unnamed generic hosts connected to a switch. If you have an unnamed generic host connected to more than one switch and you want to hide the generic element, you must repeat the following steps for each switch connected to the generic host. You can hide all unnamed generic hosts at once by using the **Hide Generic Hosts for All Switches** feature. See ["Hiding Generic Hosts for All Switches"](#) on page 203 for more information.

NOTE: This feature only pertains to unnamed generic hosts. If you name a generic host, you cannot use this feature to hide the named host. The hide generic element feature also does not work for grouped unnamed generic hosts and missing elements. To learn how to give a custom name to an unnamed generic host, see ["Assigning a Custom Name"](#) on page 221. To learn more about groups, see the topic, ["About Groups"](#) on page 197.

To hide generic hosts connected to a switch:

1. Right-click the switch.
2. Select **Discovered Element > Hide Generic Hosts for All Switches** from the drop-down menu.
A "+" icon is added to icon of the switch you right-clicked to indicate it has hidden generic hosts.

Expanding Generic Hosts for a Switch

NOTE: This feature only pertains to unnamed generic hosts. If you name a generic host, you cannot use this feature to hide the named host. The hide generic element feature also does not work for grouped unnamed generic hosts and missing elements. To learn how to give a custom name to an unnamed generic host, see ["Assigning a Custom Name"](#) on page 221. To learn more about groups, see the topic, ["About Groups"](#) on page 197.

Use the **Expand Generic Hosts for the Switch** feature to display hidden generic hosts. The hidden generic hosts are displayed in the upper-right corner of the topology if you select **Discovered Element > Expand Generic Hosts for the Switch** after you have saved the topology.

To display hidden generic hosts connected to a switch:

1. Right-click a switch with a "+" icon. This "+" icon indicates the switch has hidden generic hosts.

2. Select **Discovered Element > Expand Generic Hosts for the Switch** from the drop-down menu.

The hidden elements for the switch appear in the upper-right corner of the topology.

Hiding Generic Hosts for All Switches

NOTE: This feature only pertains to unnamed generic hosts. If you name a generic host, you cannot use this feature to hide the named host. The hide generic element feature also does not work for grouped unnamed generic hosts and missing elements. To learn how to give a custom name to an unnamed generic host, see ["Assigning a Custom Name"](#) on page 221. To learn more about groups, see the topic, ["About Groups"](#) on page 197.

You can hide all unnamed generic hosts and unnamed generic Cisco switches by doing the following:

1. Right-click a switch.
2. Select **Discovered Element > Hide Generic Hosts for All Switches** from the drop-down menu.

All unnamed generic hosts are hidden. A "+" icon is added to icon of the switch you right-clicked to indicate it has hidden generic hosts.

Expanding Generic Hosts for All Switches

NOTE: This feature only pertains to unnamed generic hosts. If you name a generic host, you cannot use this feature to hide the named host. The hide generic element feature also does not work for grouped unnamed generic hosts and missing elements. To learn how to give a custom name to an unnamed generic host, see ["Assigning a Custom Name"](#) on page 221. To learn more about groups, see the topic, ["About Groups"](#) on page 197.

Use the **Expand Generic Hosts for All Switches** feature to display hidden generic hosts. The hidden generic hosts are displayed in the upper-right corner of the topology if you select **Discovered Element > Expand Generic Hosts for All Switches** after you have saved the topology.

To display hidden generic hosts for a domain:

1. Right-click a switch with a "+" icon. This "+" icon indicates the switch has hidden generic hosts.
2. Select **Discovered Element > Expand Generic Hosts for All Switches** from the drop-down menu.

The hidden elements for the domain appear in the topology.

Setting Up Custom Commands

This section describes the following:

- ["About Custom Commands"](#) on page 204
- ["Adding a Custom Command"](#) on page 204

- ["Editing a Custom Command"](#) on page 207
- ["Deleting a Custom Command"](#) on page 207
- ["Software Environment Variables for Scripting"](#) on page 207
- ["Using the Remote Console"](#) on page 210

About Custom Commands

Custom commands let you run a command you created on the management server. The command could point to an executable or a script that does not use the graphical user interface. For example, assume you have already created a script that backs up a storage system. You could run that script from System Manager.

You can also use environment variables in your scripts. For example, you could use the variables to obtain information about a host, such as its total physical memory and the number of processors.

Keep in mind the following:

- The custom command always runs on the management server unless you are running the telnet utility. You can obtain information about the element on which you right-clicked by using the software's environment variables. See ["Software Environment Variables for Scripting"](#) on page 207.
- Run scripts at your own risk. The management server lets you run any script including those that can disable the management server.
- Custom commands only supports executables and scripts that do not use the graphical user interface.
- (Management Servers on Windows Only) If you leave the Remote Console (`cmd /k`) open after running a script, users can traverse the directory structure of the management server.
- If you want a Perl script to run as a custom command on Sun Solaris, you must prefix the script with the Perl executable, for example, `perl myscript.pl`, where `myscript.pl` is the script you want to run. A best practice is to prefix the script with the path to Perl and the Perl executable, for example: `perl/bin/perl myscript.pl`, where `perl/bin/` is the directory containing the Perl executable, `perl` is the executable and `myscript.pl` is the script you want to run.
- If you want a Perl script to run as a custom command on Microsoft Windows, you must prefix the script name with the complete path to Perl. The management server already has a directory containing the perl executable inside the folder, `%JBOS4_DIST%\server\appiq\remoteScripts\perl\bin`. You would prefix the script name as follows:
`.\perl\bin\perl myscript.pl`
 where `.\perl\bin\` is the directory containing the Perl executable in the RemoteScripts directory, `perl` is the executable and `myscript.pl` is the script you want to run.

Adding a Custom Command

Use the custom command functionality in the software to run a command you created on the management server. The command could point to an executable or a script. For example, assume you have already created a script that backs up a storage system. You could run that script from System Manager, as described in the following steps.

Keep in mind the following:

- The custom command always runs on the management server unless you are running the telnet utility. You can obtain information about the element on which you right-clicked by using the software's environment variables. See the topic, "[Software Environment Variables for Scripting](#)" on page 207.
- Run scripts at your own risk. The management server lets you run any script including those that can disable the management server.
- (Management Servers on Windows Only) If you leave the Remote Console (cmd /k) open after running a script, users can traverse the directory structure of the management server.
- If you want a Perl script to run as a custom command on Sun Solaris, you must prefix the script name with the Perl executable, for example, `perl myscript.pl`, where `myscript.pl` is the script you want to run. A best practice is to prefix the script name with the path to Perl and the Perl executable, for example: `perl/bin/perl myscript.pl`, where `perl/bin/` is the directory containing the Perl executable, `perl` is the executable and `myscript.pl` is the script you want to run.
- If you want a Perl script to run as a custom command on Microsoft Windows, you must prefix the script name with the complete path to perl. The management server already has a directory containing the perl executable inside the folder, `%JBoss4_DIST%\server\appiq\remoteScripts\perl\bin`. You would prefix the script name as follows:
`.\perl\bin\perl myscript.pl`
where `.\perl\bin\` is the directory containing the Perl executable in the RemoteScripts directory, `perl` is the executable and `myscript.pl` is the script you want to run.

To add a custom command:

1. Right-click an element in System Manager.
2. Select **Set Up Custom Commands** from the drop-down menu.
3. (Optional) If you plan to use a command to activate a file, such as a script, the file must be uploaded to the management server, as described below:
 - a. In the Custom Command Setup window, click the **Browse** button to find the file containing the custom command.
 - b. Click the **Open** button, then click the **Upload to server** button.
The file is saved on the management server.
4. Click the **Add Command** button in the upper-right corner of the window.
5. In the Add Custom Command window, type a name for the command in the **Name** field, for example, **backup command**.
6. In the **Description** field, type a description, for example, **This command activates a script that backs up an element**.
7. In the **Command Line** field type a command, which will run on the management server.
This could be a command required to start a script, for example the following:
`myscript.bat`

The remote console automatically becomes inactive once the command finishes. (Windows Only) If you want the remote console to stay open, prefix the command with the following:

```
cmd /k
```

For example:

```
cmd /k dir
```

IMPORTANT: When the command in the above example runs, it leaves the remote console window open, so you can enter additional commands. Users can traverse the directory structure of the management server when the remote console is left open.

The file is appended to the command line. (Optional) If you plan to use a file in the command, select the file from the **Files** drop-down menu. Then, click the **Append To Command Line** button.

If the file is missing, repeat step 3.

8. Select one of the following options to determine which elements you want the command to be visible. For example if you select the **All Elements** option, the command is visible in the drop-down menu when you right-click any element.
 - **Name of the Element** - Select the name of the element if you want the command to be visible in the drop-down menu when you right-click only this element.
 - **All Elements** - Select this option if you want the command to be visible in the drop-down menu when any element is right-clicked.
 - **Selected element types and filter criteria** - Select this option if you want to narrow the filtering criteria for an element type. For example, you could specify that the command is only in the drop-down menu when a Brocade switch is right-clicked. The options are as follows:
 - **Applications** - (Optional) If you want the command to be visible in the drop-down menu when a particular application is right-clicked, type the name of the product in the **Product Name** field. To make sure you type the correct product name, type the product name displayed in the Product Name field on the Properties tab, accessible by double-clicking the application in System Manager and clicking the Properties tab.
 - **Hosts** - (Optional) If you want the command to be visible in the drop-down menu when a particular host is right-clicked, type the name of the operating system in the **OS Name** field. To make sure you type the correct operating system, type the operating system displayed in the Target Operating System field on the Properties tab, accessible by double-clicking the host in System Manager and clicking the Properties tab.
 - **Switches** - (Optional) If you want the command to be visible in the drop-down menu when a switch from a particular vendor is right-clicked, type the name of the vendor in the **Vendor Name** field. To make sure you type the correct vendor name, type the vendor name displayed on the Properties tab, accessible by double-clicking the switch in System Manager and clicking the Properties tab.
 - **Storage Systems** - (Optional) If you want the command to be visible in the drop-down menu when a storage system from a particular vendor is right-clicked, type the name of the vendor in the **Vendor Name** field. To make sure you type the correct vendor name,

type the vendor name displayed on the Properties tab, accessible by double-clicking the storage system in System Manager and clicking the Properties tab.

9. Click **OK**.

10. When you want to run a command, do the following:


- a. Right-click an element.
- b. Select the command from the drop-down menu.

A remote console displays the result of the command.

You can stop a command by clicking the **Stop** button in the remote console. Once a command has been executed the console becomes inactive. The software assumes you are in the `[Install_Dir]\JBossandJetty\server\appiq\remotescripts` directory on the management server when the script is executed.

Editing a Custom Command


To edit a custom command :

1. Right-click an element in System Manager.
2. Select **Set Up Custom Commands** from the drop-down menu.
3. Click the  button corresponding to the custom command you want to edit.
4. Make the appropriate changes in the **Edit Custom Command** window.
5. Click **OK**.

The custom command is modified.

Deleting a Custom Command

To delete a custom command:

1. Right-click an element in System Manager.
2. Select **Set Up Custom Commands** from the drop-down menu.
3. Click the  button corresponding to the custom command you want to delete.

The custom command is deleted.

Software Environment Variables for Scripting

Keep in mind the following:

- Run scripts at your own risk. The management server lets you run any script including those that can disable the management server.
- (Management Servers on Windows Only) If you leave the Remote Console (`cmd /k`) open after running a script, users can traverse the directory structure of the management server.

The software provides the following environment variables for you to put in your scripts. For example, assume you have a script that backups a host. You could use the variables to obtain information about the host.

The software gathers information about the element you right-clicked. For example, if you use the variable, `APPIQ_ELEMENT_ID`, the management server obtains the element ID of the element you right-clicked.

The following table lists the variables that can be used to gather information for all elements.

Table 30 Variables for All Elements*

Variable	Value
<code>APPIQ_ELEMENT_ID</code>	The identifier of an element.
<code>APPIQ_ELEMENT_NAME</code>	The name of the element.
<code>APPIQ_ELEMENT_STATUS</code>	The status of the element, such as the following: <ul style="list-style-type: none"> • Managed • Generic • Missing • Virtual Application • Asset
<code>APPIQ_ELEMENT_DESCRIPTION</code>	The description for the element.
<code>APPIQ_ELEMENT_MANAGEMENT_IP_ADDRESS</code>	Contains the IP address of the first access point used to discover the element.
<code>APPIQ_ELEMENT_VENDOR</code>	The vendor for the element.
<code>APPIQ_ELEMENT_TYPE_NAME</code>	The type of element, for example, switch, application, or host.

*If an application resides on the host, the variables in this table provide information about the application. See the last table for variables that return information about the host.

The following table lists variables that can be used to gather information for only storage systems, switches, and hosts.

Table 31 Variables for Only Storage Systems, Switches, and Hosts*

Variable	Value
<code>APPIQ_ELEMENT_IP_ADDRESS</code>	The IP address of the element.
<code>APPIQ_ELEMENT_DNS_NAME</code>	The DNS name of the element.
<code>APPIQ_ELEMENT_MODEL</code>	The model of the element.
<code>APPIQ_ELEMENT_VERSION</code>	The version of the element.

*If an application resides on the host, the variables in this table provide information about the application. See the last table for variables that return information about the host.

The following table lists variables that can be used to gather information for only switches.

Table 32 Variables for Only Switches

Variable	Value
APPIQ_ELEMENT_SWITCH_ID	The identifier for the switch.
APPIQ_ELEMENT_IP_GATEWAY	The IP gateway of the switch.
APPIQ_ELEMENT_IP_NETWORK_MASK	The IP network mask for the switch.
APPIQ_ELEMENT_SWITCH_STATUS	The status of the switch.
APPIQ_ELEMENT_DOMAIN_ID	The domain identifier of the switch.

The following table lists variables that can be used to gather information for only hosts.

Table 33 Variables for Only Hosts*

Variable	Value
APPIQ_ELEMENT_OPERATING_SYSTEM	The operating system of the host.
APPIQ_ELEMENT_NUMBER_OF_PROCESSORS	The number of processors used by the host.
APPIQ_ELEMENT_TOTAL_PHYSICAL_MEMORY	The total physical memory of the host.
APPIQ_ELEMENT_DOMAIN	The domain of the host.

*If an application resides on the host, the variables in this table provide information about the application. See the last table for variables that return information about the host.

The following table lists variables that can be used to gather information for applications and hosts. Use the variables with the "APPIQ_HOST" prefix when you are using variables from the first table to gather information about the application. For example, if you are running a script containing APPIQ_ELEMENT_STATUS on a host, it would obtain information about the status of the application. You would need to run APPIQ_HOST_STATUS to obtain information about the status of the host on which the application resides.

Table 34 Variables for Only Applications

Variable	Value
APPIQ_ELEMENT_PRODUCT_NAME	The name of the application.
APPIQ_HOST_NAME	The name of the host on which the application resides
APPIQ_HOST_ID	The identifier of a host on which the application resides.
APPIQ_HOST_STATUS	The status of the host on which the application resides.
APPIQ_HOST_DESCRIPTION	The description of the host on which the application resides.
APPIQ_HOST_VENDOR	The vendor of the host on which the application resides.
APPIQ_HOST_TYPE_NAME	The name of the host on which the application resides.
APPIQ_HOST_IP_ADDRESS	The IP address of the host on which the application resides.
APPIQ_HOST_DNS_NAME	The DNS name of the host on which the application resides.
APPIQ_HOST_MODEL	The model of the host on which the application resides.
APPIQ_HOST_VERSION	The version of the host on which the application resides.
APPIQ_HOST_OPERATING_SYSTEM	The operating system of the host on which the application resides.
APPIQ_HOST_NUMBER_OF_PROCESSORS	The number of processors on the host on which the application resides.
APPIQ_HOST_TOTAL_PHYSICAL_MEMORY	The total physical memory of the host on which the application resides.
APPIQ_HOST_DOMAIN	The domain of the host on which the application resides.

Using the Remote Console

This section describes the following:

- ["About the Remote Console"](#) on page 211

- ["Keeping the Remote Console Active"](#) on page 211
- ["About the Buttons on the Remote Console"](#) on page 212
- ["About the Drop-Down Menu Options"](#) on page 212
- ["Copying Text from the Remote Console"](#) on page 213

About the Remote Console

Whenever you run a custom command on the management server, the remote console appears. The remote console displays the result of a custom command. For example, you can use the remote console to start a remote command prompt on the management server.




Figure 39 Remote Console

Keeping the Remote Console Active

NOTE: This section is for only management servers running on Microsoft Windows.

Keep in mind the following:

- The remote console become inactive when the custom command finishes its execution. To use the menus and buttons in the remote console, the remote console must be kept active.
- If you leave the Remote Console (`cmd /k`) open after running a script, users can traverse the directory structure of the management server.

To keep the remote console window active, create a remote command prompt:

1. Right-click an element in System Manager.

2. Select **Set Up Custom Commands** from the drop-down menu.
3. Click the **Add Command** button in the upper-right corner of the window.
4. In the **Add Custom Command** window, type a name for the command in the **Name** field, for example, **prompt**.
5. In the **Description** field, type a description, for example, **Accesses the remote console**.
6. In the **Command Line** field type the following command, which will run on the management server:

```
cmd /k
```
7. Select the **All elements** option.
8. Click **OK**.
9. When you want to run the remote command prompt, do the following:
 - a. Right-click an element from which you want to obtain information.
 - b. Select the command from the drop-down menu.

The software displays the remote console on the management server.

10. To enter a command in the remote console, type the command in the **Commands** field. Then, press ENTER.

You can stop a command by clicking the **Stop** button in the remote console.

Keep in mind the following:

- You can quickly access information about the element you right-clicked by typing the following at the command prompt:

```
set appiq
```

- The software ships with a utility called plink. To view the commands for plink, type the following in the **Commands** field and then press ENTER:

```
plink
```

About the Buttons on the Remote Console

The remote console provides the **Stop** and **Help** buttons, as described in the following table.

Table 35 Buttons on the Remote Console

Button	Description
Stop	Stops a command. Once a command has been executed the console becomes inactive.

Table 35 Buttons on the Remote Console (continued)

Button	Description
Help	<p>Provides the following information about the remote console:</p> <ul style="list-style-type: none">• Clearing the remote console - Type <code>CLS</code> in the Commands field of the remote console.• Copying text to the Commands field - Place the cursor at the end of the line in the window below the Commands field, as shown in the following figure. Then, press ENTER. The command is copied to the Commands field. <p>Note: If you are viewing the remote console on Microsoft Windows, you can copy text by using CTRL + C, then use CTRL + P to paste it.</p>

About the Drop-Down Menu Options

The remote console also provides the following drop-down menu options.

Table 36 Drop-Down Menu Options

Option	Description
CRLF	<p>(Default setting) Provides a carriage return and a linefeed.</p> <p>Important: Do not use this option when you are using telnet to access another computer. You must select the CR option after you enter a user name to be able to enter a password. If you leave the setting at CRLF, the software enters a carriage return and a line feed when you click OK. As a result, no value is entered for the password.</p>
CR	Provides a carriage return.
LF	Provides a linefeed.

Copying Text from the Remote Console

To copy text from the remote console:

1. Select the text in the remote console.
2. Right-click the top frame in the remote console.
3. Select **Copy** from the drop-down menu.

The text is stored in the buffer of your computer to be pasted elsewhere.

Using External Tools

This section describes the following:

- ["About the External Tools Feature"](#) on page 213

- **Set Up External Tools** - Lets you add a URL for accessing management software, such as Hitachi HiCommand Device Manager and EMC ControlCenter™ Navisphere. See "[Setting up External Tools](#)" on page 214.

About the External Tools Feature

The management server ships with an external tools feature that lets you:

- **Browse the Element** - Access a host or a switch through its main Web page. The software assumes the host or switch has a Web page at `http://hsIPAddress`, where `hsIPAddress` is the IP address of the host or switch. To access the main Web page of the host or switch, right-click the element in System Manager and select **External Tools > Browse to 192.168.1.2**, where 192.168.1.2 is the IP address of the host or switch.
- **Telnet to the Element** - Access a host or a switch through the telnet utility. Telnet must be already enabled on the element. The command uses `telnet://hsIPAddress`, where `hsIPAddress` is the IP address of the host or switch. To telnet to a host or switch, right-click the element in System Manager and select **External Tools > Telnet to 192.168.1.2**, where 192.168.1.2 is the IP address of the host or switch.
- **Set Up External Tools** - Lets you add a URL for accessing management software, such as Hitachi HiCommand Device Manager and EMC ControlCenter™ Navisphere. See "[Setting up External Tools](#)" on page 214 for more information.
- **Access the Management Tool for the Storage System** - In some instances, the management tool for the storage system is accessible from this menu. For example, HiCommand for HDS storage systems and Command View for HP XP storage systems are accessible from the External Tools menu.

Setting up External Tools


You can add URLs used for accessing external tools used for managing an element, such as Hitachi HiCommand Device Manager and EMC ControlCenter™ Navisphere for storage systems.

NOTE: When you add a URL, it only applies to the element you originally right-clicked.

To add a URL for accessing external tools:

1. Access System Manager.
2. In the right pane, right-click the element and select **External Tools > Set Up External Tools**.
3. Click the **Add New Management URL** button.
4. In the **Description** field, type the name of the product you plan to access.
5. In the **URL** field, type the URL that is used to access the product.
6. Click **OK**.

When you right-click the element and select **External Tools**, the external tool is listed.

To delete the URL for an external tool, click the corresponding  button in the External Tools window.

About the Navigation Tab

The Navigation tab not only provides information about an element, but it also illustrates how the element relates to other elements in its path. For example, the Navigation page displays logical and physical components, such as ports, zone sets, zones and zone aliases. It also displays the dependencies for switches, as shown in the following figure.

Keep in mind the following:

- When McDATA or Connectrix switches are discovered through a proxy by using SNMP, you cannot view or perform any provisioning operations for those switches. For example, you cannot view zone sets, zones, and/or zone alias.
- When McDATA or Connectrix switches are discovered by their IP address by using SNMP, you can only view the active zone set and its members. You cannot create, modify, and/or delete zone sets or its members.

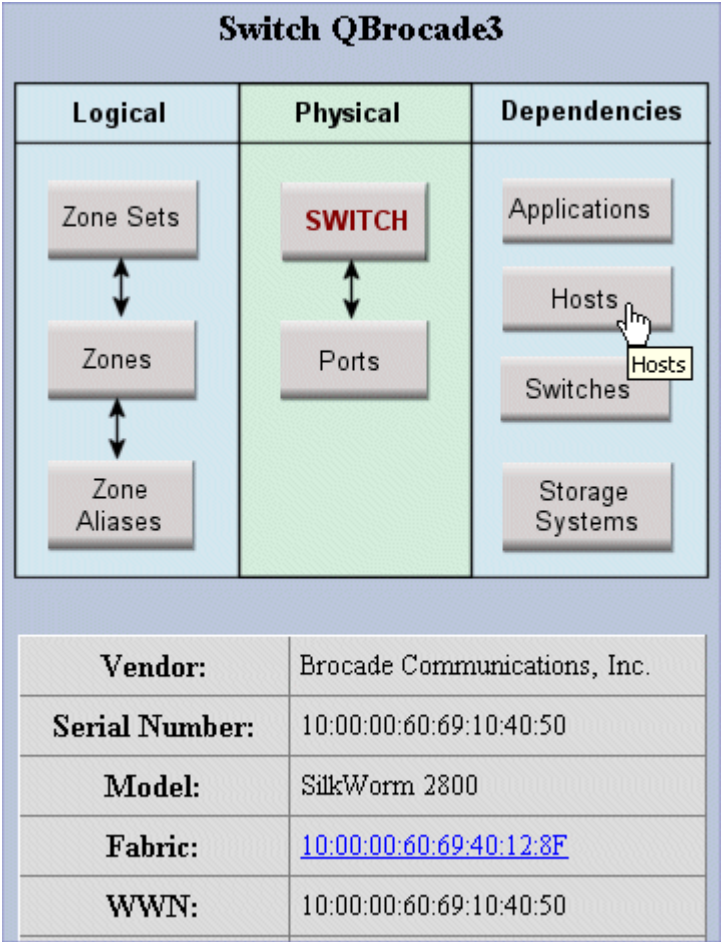


Figure 40 Obtaining Information About a Host

You can learn more about a component, by clicking it in the Navigation page. For example, assume you brought up the Navigation page shown in the previous figure, and you want to learn which hosts are dependent. You would click Hosts in the page. You are shown information about the dependent hosts, as in the following figure:

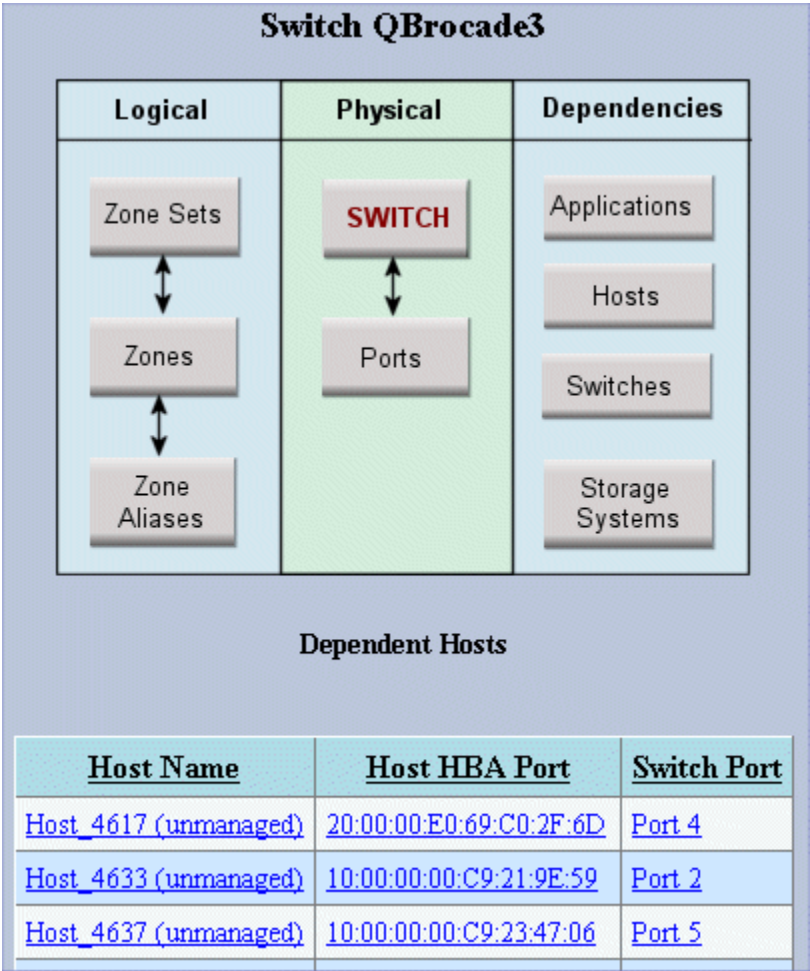


Figure 41 Details of a Host Connected to a Switch

The following table provides an overview of the information presented for each type of element.

Table 37 Information Available from the Navigation Page

Element	Dependencies	Front Physical	Back Physical	Logical	Physical
Applications	✓				

Table 37 Information Available from the Navigation Page (continued)

Element	Dependencies	Front Physical	Back Physical	Logical	Physical
*Hosts	✓				
Switches	✓			✓	✓
Storage Systems	✓	✓	✓	✓	

*The management server displays cxfs for SGI IRIX computers if it detects cxfs on the cluster. On individual IRIX computers "cxfs" is not displayed when you type the following at the command prompt:

```
df -k
```

Data may be missing from the Navigation tab for a McDATA or a Connectrix switch if the switch was discovered by using SNMP and one of the following techniques:

- Its IP address
- Enterprise Fabric Connectivity (EFC) Manager or EMC Enterprise Connectrix Manager (ECM)

If the McDATA and/or Connectrix switches are discovered by their IP address (SNMP connection), the following fields are empty:

- IP Gateway
- Switch ID
- FC Net Address
- FC Net Mask

If the McDATA and/or Connectrix switches are discovered by SNMP and by Enterprise Fabric Connectivity (EFC) Manager or by EMC Enterprise Connectrix Manager (ECM), the following fields are empty:

- IP Gateway
- Switch ID
- DNS Name
- IP Address
- IP Net Mask
- FC Net Address
- FC Net Mask

To learn how to access the Navigation tab, see "[Accessing the Navigation Tab](#)" on page 219 for more information.

Finding the Status of a Port on a Switch

The management server can detect the status of a switch. This can be especially useful if a port is being a problem. To find the status of a port on a switch:

1. Access System Manager as described in "[Accessing System Manager](#)" on page 163.
2. Double-click a switch in the right pane.
3. Click the **Ports** button under the Physical column in the Navigation tab, as shown in the following figure.

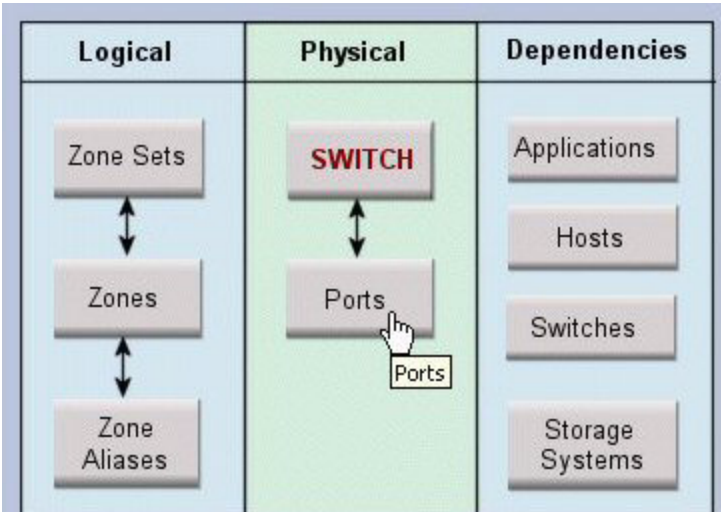


Figure 42 Finding the Status of a Port

4. Under the Name column in the Ports table, click the port that you want to obtain the status.
5. On the Properties page, the status of the port is displayed in the right column.
The status of the port can be online, off line or unknown.

Table 38 Port Status Definitions

Status	Definition
Online	Port is physically installed with node connections
Offline	Port is physically installed, but without node connections Brocade switches also display this status if the port is not physically installed (Gigabit Interface Converter (GIBIC) is not installed).
*Not Installed (McDATA SWAPI connections only)	Port is not physically installed (Gigabit Interface Converter (GIBIC) is not installed).

Table 38 Port Status Definitions (continued)

Status	Definition
*Unknown (McDATA SNMP connections only)	Port is not installed.

*An ES4500 switch displays its status differently when a port is not installed:

- **SWAPI connection** - Unknown status
- **SNMP connection** - Offline

Accessing the Navigation Tab

To access the **Navigation** tab:

1. Access the management server.
2. To access the **Navigation** tab, do one of the following:
 - Click an element in Application Viewer.
 - Double-click an element in Capacity Manager, Performance Manager or System Manager.
3. Click the **Navigation** tab.

Viewing Element Properties

This section describes the following:

- [“About the Properties Tab”](#) on page 219
- [“Accessing the Properties Tab”](#) on page 220
- [“Viewing Fabric Properties”](#) on page 220
- [“Assigning a Custom Name”](#) on page 221
- [“Adding a Generic Tape Library”](#) on page 222

About the Properties Tab

The Properties tab provides detailed information about an element. Since the information obtained from each type of element varies, the Properties tab displays information only relevant for that type of element. For example, the Properties tab for fabrics lists the zones, zone sets, switches, and zone aliases, as compared to the Properties tab for a host, which lists the processors, cards, applications, and storage volumes the host uses.

See [“Viewing Fabric Properties”](#) on page 220 for more information about the Properties tab for fabrics.

The Properties tab usually provides the following, although this list does vary from element to element:

- ***Assign a Custom Name** - To make it easier to identify the element, assign the element a Custom Name. See the topic, [“Assigning a Custom Name”](#) on page 221. This option is not available to all elements.
- ***IP Address (Generic Hosts Only)** - Type an IP address for a generic host.

- ***DNS Name (Generic Hosts Only)** - Type a DNS name for a generic host.
- ***Version (Generic Hosts Only)** - Type a version number for the generic host.
- ***Operating System (Generic Hosts Only)** - Type an operating system for a generic host.
- **Vendor** - The vendor name.
- **Tape Library option** - Select this option if the element is a tape library. See ["Adding a Generic Tape Library"](#) on page 222.
- **View element properties** - The element properties viewed on the type of device. The properties provided vary according to the element. The following information is usually provided:
 - **Record Creation** - The first time the software contacted this element.
 - **Discovery Status** - The status of the discovery of the element, for example "Contacted."
 - **Vendor** - The name of the vendor.
 - **IP Address** - The IP address of the element.
 - **DNS Name** - The element's domain name system (DNS) name.
 - **Provider Name** - The name of the provider.
 - **Model** - The model of the element.
- **Update Element Data** - To update the displayed properties, click the **Update Element Data** button at the bottom of the screen. The management server gathers new and changed details from the element and then redraws the topology with the updated information.

Important:

- Do not update element data during "Discovery Data Collection." You can determine if the management server is getting the topology or all element details by looking at the label near the status button.
- The Update Element Data functionality does not detect element components that have been removed, such as ports and LUNs. For example, assume you removed several LUNs from an array. If you perform **Update Element Data** for the storage system, the LUNs still appear in the user interface. You must perform "Discovery Data Collection" for the deleted LUNs to be removed from the user interface. See the topic, ["Updating the Database with Element Changes"](#) on page 45 for more information about "Discovery Data Collection."

*Click the **Save** button for your changes to take effect.

Accessing the Properties Tab

To access the **Properties** tab:

1. Access the management server.
2. To access the **Properties** tab for an element, do one of the following:
 - Click an element, except a file server, in Application Viewer.
 - Double-click an element in System Manager.
3. To access the **Properties** tab for a fabric:
 - a. In System Manager, click the **List** tab.
 - b. Right-click a fabric name in the **List** tab, for example 100000606930260d.

- c. Select **Go to Properties** from the right-click menu.
4. (Not applicable to fabrics)Click the **Properties** tab.

Viewing Fabric Properties

You can view the following properties of a fabric:

- **Assign a Custom Name** - To make it easier to identify the element, assign the element a Custom Name. See "[Assigning a Custom Name](#)" on page 221. This option is not available to all elements.
- **Vendor** - The vendor name.
- **Created** - The first time the software contacted this element.
- **Discovery Status** - The status of the discovery of the element, for example "Contacted."
- **Install Date** - Not applicable
- **Name Detected** - The name of the fabric detected.
- **OID** - Not applicable
- **Description** - The status of the discovery of the element, for example "Contacted."
- **WWN** - The Worldwide Name of the fabric.
- **Zones** - The zones in the fabric. To learn more about a zone, click its link.
- **Zone Sets** - The zone sets in the fabric. To learn more about a zone set, click its link.
- **Switches** - The switches in the fabric. To learn more about a switch, click its link.
- **Zone Aliases** - The zone aliases in the fabric.
- **Update Element Data** - To update the displayed properties, click the **Update Element Data** button at the bottom of the screen. The management server gathers new and changed details from the element and then redraws the topology with the updated information.

Keep in mind the following:

- Do not update element data during "Discovery Data Collection." You can determine if the management server is getting the topology or all element details by looking at the label near the status button.
- The Update Element Data functionality does not detect element components that have been removed, such as ports and LUNs. For example, assume you removed several LUNs from an array. If you perform **Update Element Data** for the storage system, the LUNs still appear in the user interface. You must perform "Discovery Data Collection" for the deleted LUNs to be removed from the user interface. See the topic, "[Updating the Database with Element Changes](#)" on page 45 for more information about "Discovery Data Collection."

To learn how to access the Properties tab for a fabric, see the topic, "[Accessing the Properties Tab](#)" on page 220.

Assigning a Custom Name

To make it easier to identify an element instance in the system, assign the instance a Custom Name. The Custom Name also appears in Chargeback.

IMPORTANT: Do not assign a custom name during "Discovery Data Collection." You can determine if the management server is getting the topology or all element details by looking at label near the status button.

NOTE: Since all users query the same database, this name is displayed to others using the software. As a result, you might want to make them aware of the name.

1. Access the **Custom Name** field by double-clicking the element in System Manager and then clicking the **Properties** tab.
2. In the **Custom Name** field, type a name.
Keep in mind the following:
 - The name must contain 1 to 64 characters.
 - The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or one of the following symbols: dollar sign (\$), caret (^), hyphen (-), an underscore (_) or a space.
 - The name is case sensitive, for example, "Element1" and "element1" are different elements.
3. Click **Save**.

Adding a Generic Tape Library

You can convert a discovered storage system to a discovered tape drive, meaning the tape drive will not be managed or monitored by the management server. The icon is a placeholder in the topology in System Manager. If you want to manage and monitor the tape library, see ["Discovering IBM Tape Libraries"](#) on page 37.

Keep in mind the following:

- The management server does not monitor the tape drive.
- Since all users query the same database, the custom name you assign to the tape library is displayed to others using the software. As a result, you might want to make them aware of the new name.
- Do not assign information to the tape drive during "Discovery Data Collection." You can determine if the management server is getting the topology or all element details by looking at label near the status button.

To assign information to the tape drive:

1. Double-click the discovered storage system in System Manager.
2. In the **Custom Name** field, type a name.
Keep in mind the following:
 - The name must contain 1 to 64 characters.
 - The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or one of the following symbols: dollar sign (\$), caret (^), hyphen (-), an underscore (_) or a space.

- The name is case sensitive, for example, "Element1" and "element1" are different elements.
3. In the **Vendor** field, type the vendor of the tape drive.
 4. Select the **Tape Library** option.
 5. Click the **Save** button.

Viewing Element Topology

This section describes the following:

- ["About the Topology Tab"](#) on page 222
- ["About the New Window Option"](#) on page 228
- ["Printing the Topology"](#) on page 229

About the Topology Tab

The Topology tab provides a graphical representation of an element's path. It displays additional information not found in System Manager, such as adapters, slots, and fibre channel ports.

For example, assume you want to view the topology of a server called QASERVER02, and it contains seven fixed local disks. If you double-click the server in System Manager and then click the **Topology** tab, you can see the path of the server. The Topology tab also displays the drives of the server's fixed local disks, as well as the adapter used to connect the server with the switch, as

shown in the following figure. According to the following figure, the server can access three storage systems: Engenio, EMC, and HDS.

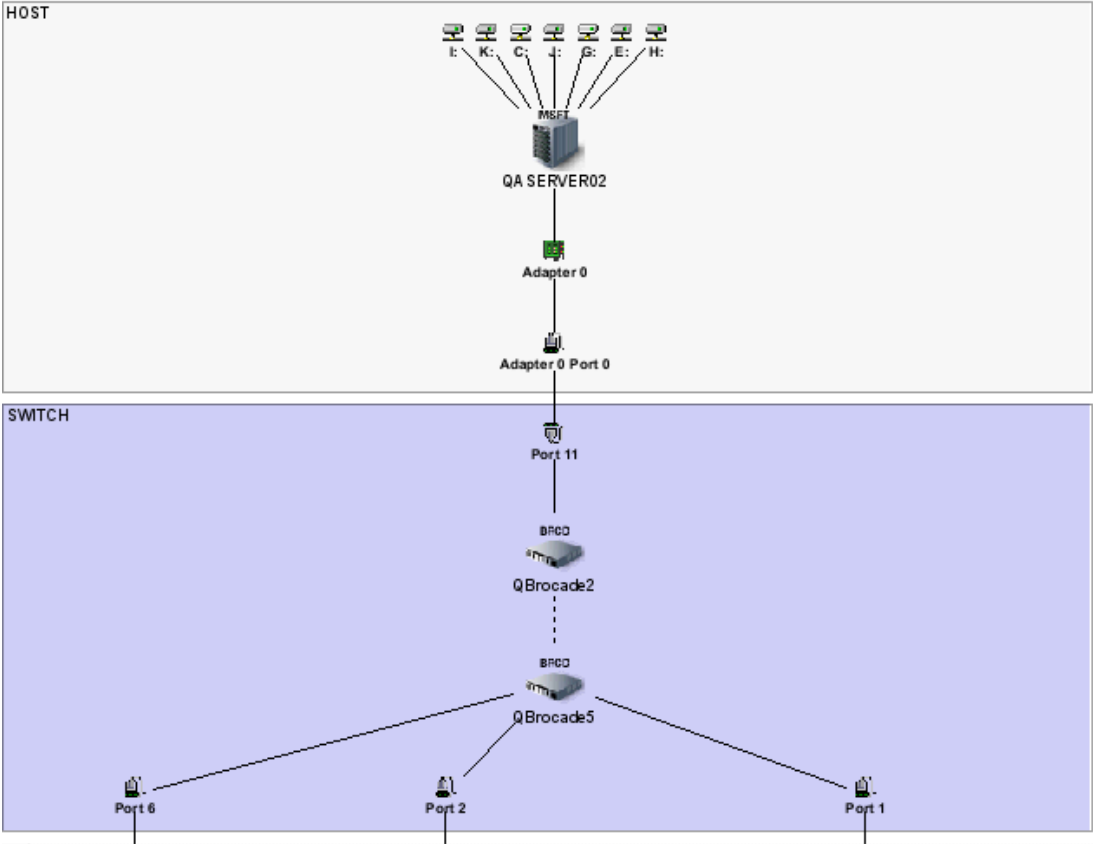


Figure 43 Topology of a Server

The topology extends the length of the screen. The second portion of the topology is provided by the following figure.

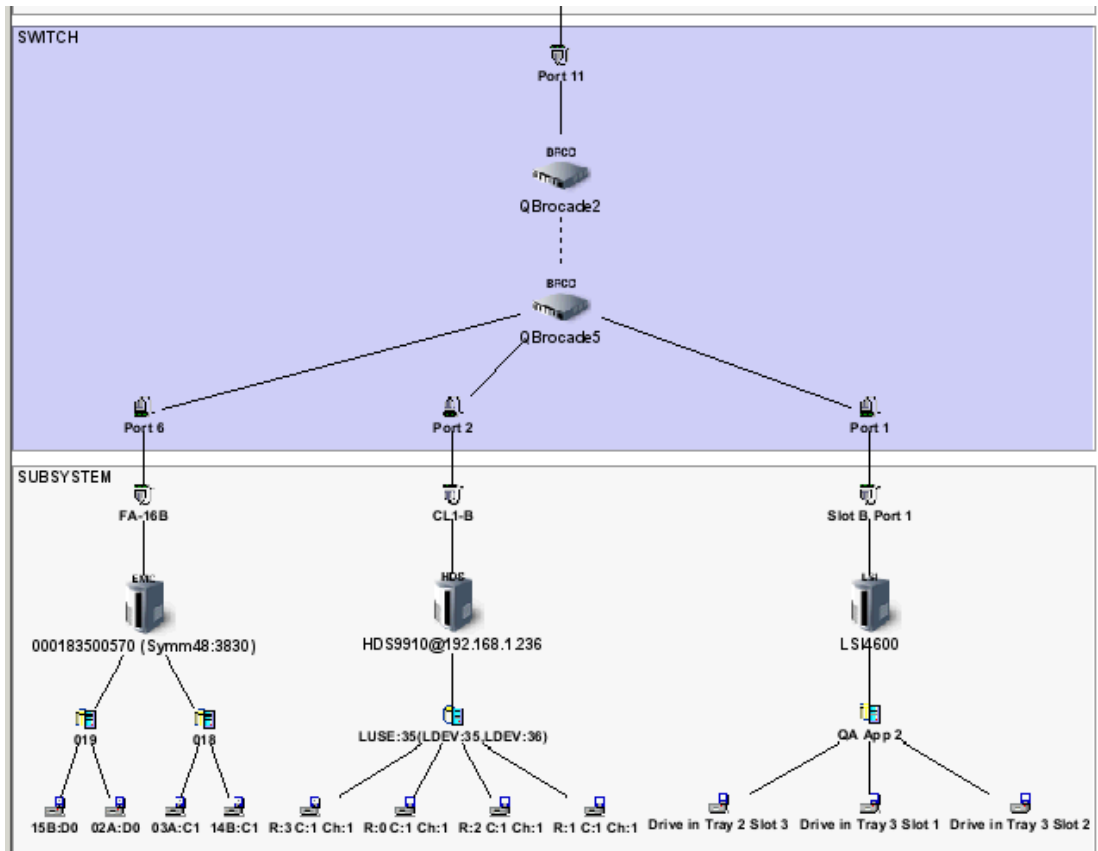


Figure 44 Topology of a Server (Continued)

Multipathing

Multipathing is the process of providing a server more than one path to a storage system, so that in the case of an emergency, the server will have continuous access to the storage system. Multipathing can be done many ways. For example, you can provide redundant switches for a host to access a storage system. Another example of multipathing is providing redundant paths from the host to the switch. To determine if your multipathing software is supported, refer to the release notes.

Keep in mind the following:

- "SANtricity" Manager Utilities" must be installed on the host running RDAC for the management server to obtain RDAC information.
- HDLM on Sun Solaris requires the storage array to be included in discovery in order to report the correct information back to the bindings page. For more information about HDLM issues, see "Troubleshooting HDLM."

- The software supports VERITAS Volume Manager without VxDMP, but VxDMP is required to do multipathing.

The following figure displays how the software detects multipathing for a server running VERITAS Volume Manager. "MP" is displayed on the path of the redundant volumes.

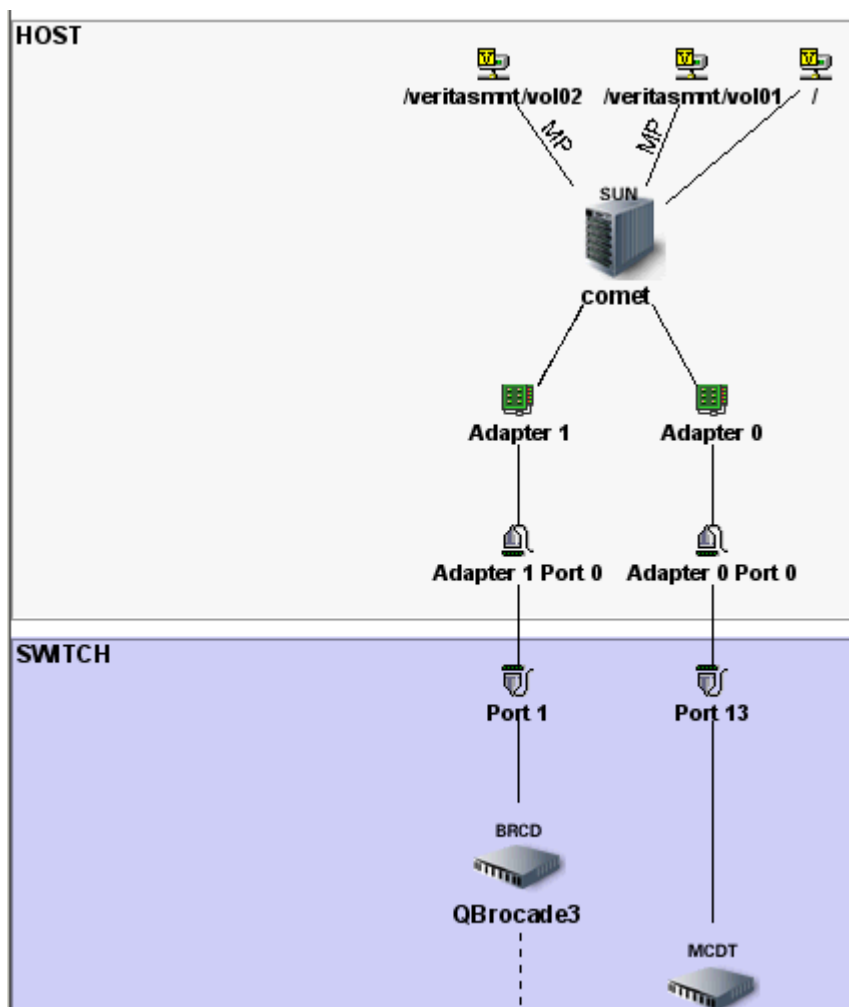


Figure 45 Multipathing Displayed in the Topology

The topology extends the length of the screen. The second portion of the topology is provided by the following figure.

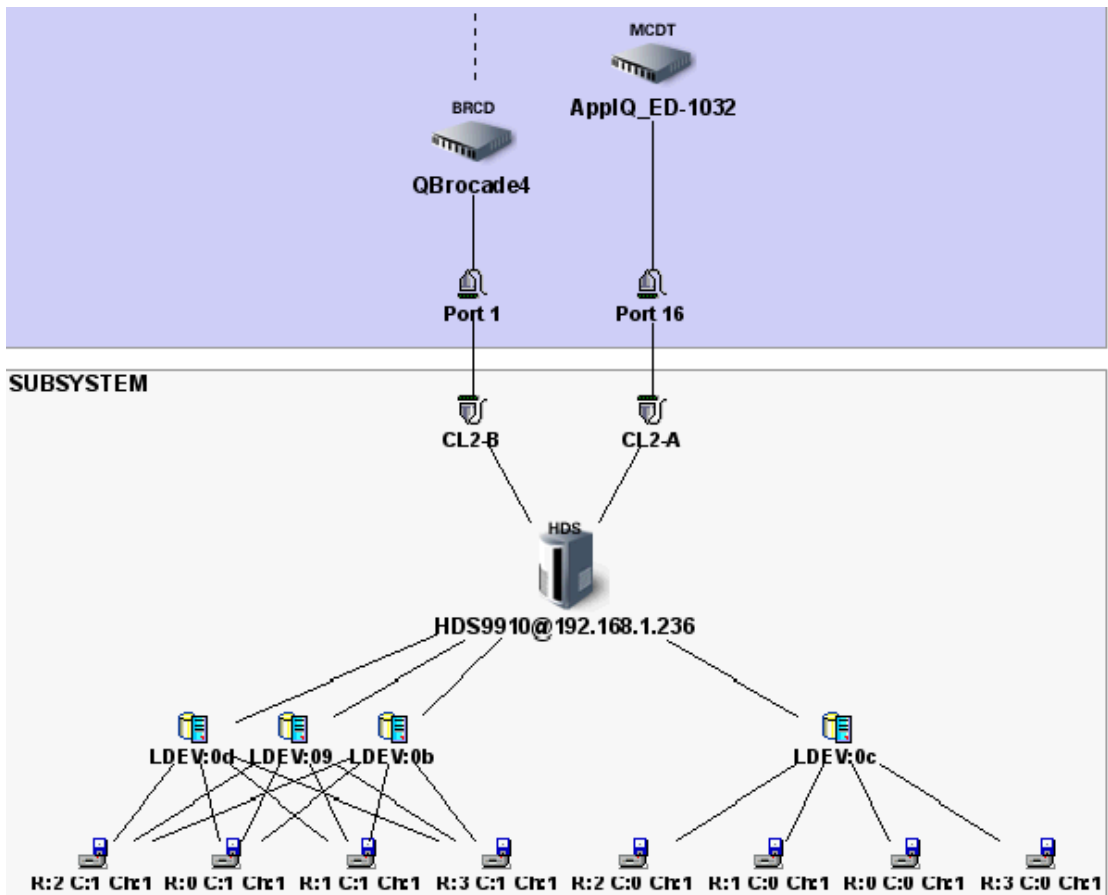




Figure 46 Multipathing Displayed in the Topology (Continued)

Keep in mind the following:

- If you do not see all of the elements in the path displayed, verify they have been discovered and details have been obtained from them. See ["Discovering Filers, Tape Libraries, Switches, and Storage Systems"](#) on page 13.
- See ["Known Device Issues"](#) on page 524 for multipathing issues regarding certain devices.
- The management server displays only the active path for an RDAC host. It also displays only the active path when PowerPath is running on a host connected to a CLARiiON storage system. The management server also does not support RDAC configurations for monitoring disk statistics.

Direct Attached Storage

To view direct attach storage, you must enable the  button. See [Table 23](#) on page 164 for more information.

Once the  button is enabled, the management server displays the link between the storage system port and the port to the host as a dotted line, as shown in the following figure:




Figure 47 Direct Attached Storage in the Topology

In this figure, Slot A Port 1 belongs to the storage system and Adapter 1 Port 0 belongs to the host. The dotted line indicates the storage system is directly attached to the host.

Filers




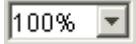






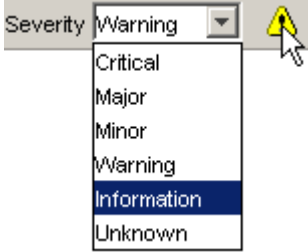
Element topology for a filer shows the connection from a host to the filer going through an IP cloud, which represents the IP network.

Accessing the Topology


To access the Topology tab:

1. Access the management server.
2. Access the Topology tab by doing one of the following:
 - Selecting an element in Application Viewer, and then clicking the **Topology** tab.
 - Double-clicking an element in Capacity Manager or System Manager, and then clicking the **Topology** tab.
 - Right-clicking an element in System Manager, and then selecting **Show Element Topology** from the drop-down menu.

Table 39 About the Toolbar in the Topology Tab

Icon	Description
	Prints the topology. See the topic, " Printing the Topology " on page 229 for more information.
	Magnifies the view
	Decreases the magnification
	Lets you set the magnification to a percentage of the default magnification
	Opens a smaller pane, which provides a global view of the topology. This lets you position the main view to a certain section of the topology. For more information, see " Using the Global View " on page 187.
	Lets you drag an element in the topology.
	Lets you move the entire topology at once. Click the Pan button () and then click any place in the topology. Drag the mouse to a new location.
	<p>Lets you find an element by name or by Worldwide Name in the topology. You can type part of the information, and the management server highlights the elements that match.</p> <p>After you populate the search field, click the  button or press ENTER.</p>
	Displays the event severity icons for the elements displayed in the topology. See " Viewing Event Status in the Topology " on page 194 for more information. Disabled for Performance Manager and Capacity Manager.

About the New Window Option

The New Window Option in System Manager lets you view several sections of the topology at once. Click the  button. A new window pops up. Use this pane to view another section of the topology.

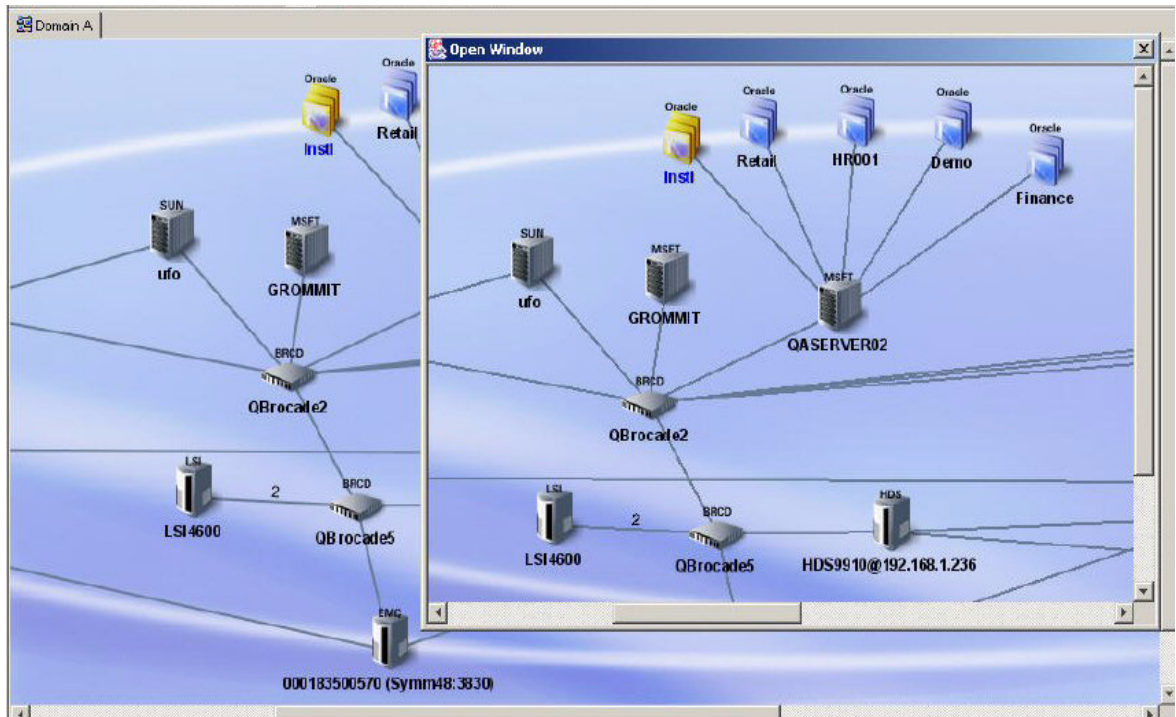




Figure 48 New Window Option

Printing the Topology

The software lets you print the topology. This option is extremely helpful when you want to show someone the layout of the network, such as in a presentation.

To print the topology:

1. Access the management server.
2. Access the Topology tab by doing one of the following:
 - Selecting an element in Application Viewer, and then clicking the **Topology** tab.
 - Double-clicking an element in Capacity Manager or System Manager, and then clicking the **Topology** tab.
3. If the topology spans more than one screen, arrange the elements so they are closer together, preferably on one screen. To move an element, click the  button and then the element you want to move. Drag the element to its new location. Moving elements closer together prevents the printout from appearing too stretched.
4. Click the  button.
5. Use the fields on the Paper tab to modify the setup of the page. When you are done, click **Apply**. If you want the default settings, click **Default**.
A preview of the printout is displayed in the right pane.

IMPORTANT: Before you change the margins, decide on a unit of measurement.

- **Paper format** - Select the paper size from the drop-down menu.
 - **Unit** - Select cm (centimeters) or inch for the margins.
 - **Paper width** - Displays the width of the paper. You can modify the measurement in this field when you select the **Custom** option in the Paper format drop-down menu.
 - **Paper height** - Displays the height of the paper. You can modify the measurement in this field when you select the **Custom** option in the Paper format drop-down menu.
 - **Top margin** - Type a measurement for the top margin.
 - **Bottom margin** - Type a measurement for the bottom margin.
 - **Left margin** - Type a measurement for the left margin.
 - **Right margin** - Type a measurement for the right margin.
 - **Orientation** - Click an orientation for the printout.
6. Click the **View Selection** tab to modify how the printout will appear on the page. You can modify the following. When you are done, click **Apply**. If you want the default settings, click **Default**.
- A preview of the printout is displayed in the right pane.

IMPORTANT: Before you change the margins, decide on a unit of measurement.

- **Start x** - Determines the horizontal placement of the printout on the page with zero being the closest to the right margin. For example, if the value is 50 for **Start x**, the printing starts at 50 inches or centimeters (depending on what you selected) from the right margin. You can also enter negative numbers. Anything more than zero expands the printout to another page.
 - **Start y** - Determines the vertical placement of the printout on the page with zero being the closest to the bottom margin. For example, if the value is 50 for **Start y**, the printing starts at 50 inches or centimeters (depending on what you selected) from the bottom. You can also enter negative numbers.
 - **Width** - Determines the width of the printout.
 - **Height** - Determines the height of the printout.
- To remove extra space around the topology, click the **Trimmed** button.
7. To change how many pages the printout will use, select one of the following. When you are done, click **Apply**. If you want the default settings, click **Default**.
- A preview of the printout is displayed in the right pane.

IMPORTANT: Before you change the margins, decide on a unit of measurement.

- **Unit** - Select cm (centimeters) or inch for the margins.

- **Position/Size** - Lets you change the position and size of the printout so that it spans several pages:
 - **Start x** - Determines the horizontal placement of the printout. Anything more than zero expands the printout to another page.
 - **Start y** - Determines the vertical placement of the printout.
 - **Width** - Determines the width of the printout. If the width entered does not fit on the page, the printout wraps around to another page.
 - **Height** - Determines the height of the printout. If the height entered does not fit on the page, the printout wraps around to another page.
 - **Resolution (pixel/unit)** - Lets you change the resolution so that the printout spans several pages.
 - **Page** - Lets you expand the printout so it prints on several pages without modifying the graphic.
8. To preview your pages, click the Preview tab. Then click the page you want to preview. The page appears in the right pane.
 9. When you are ready to print, click **Print**.
 10. Click **Close**.

NOTE: To revert back to all of the original settings, click the **Default** button next to the **Print** button.

Creating a Virtual Application

The management server lets you keep track of unsupported applications. For example, assume your company has created an internal application, and you want to be able to use the software to keep track of that application. You can create a virtual application for that product. A virtual application is a placeholder you create for an application.

Once you create the virtual application, it will appear connected to a host in your topology.

1. Access System Manager by clicking the **System Manager** button in the left pane.
2. Right-click the host that contains the application you want to monitor.
If the host is not in the topology, verify you have discovered the element and obtained element details. See "[Discovering Filers, Tape Libraries, Switches, and Storage Systems](#)" on page 13.
3. Select **Add Virtual Application** from the drop-down menu.
4. Type the following information for the virtual application.
 - **Name**
 - **Product**
 - **Description**
 - **Vendor**
 - **Version**
5. Click **Next**.

6. Select a storage volume containing the application for which you are creating the virtual application.

NOTE: You can view the properties of a volume by clicking its link.

7. If applicable, choose a disk partition by clicking the **Disk Partitions** tab or the **Next** button and then selecting a disk partition.
8. Click **Finish**.

About the Provisioning Tab

The provisioning tab provides different functionality depending on the type of element you double-clicked in System Manager or clicked in the Provisioning pages. You can also access the provisioning table by right-clicking a fabric, selecting the **Go to Properties** option, and clicking the Provisioning tab.

If you selected a switch or a fabric, you are shown zone provisioning tools that let you manage zones, zone aliases, and zone sets. These tools provide a wide range of functionality, such as the following:

- ["Creating Zone Aliases"](#) on page 245
- ["Creating a Zone in a Fabric"](#) on page 247
- ["Creating a Zone Set"](#) on page 249
- ["Activating a Zone Set"](#) on page 253

For more information about setting up zones, see ["SAN Zoning Overview"](#) on page 240.

If you double-clicked a storage system, you are shown storage provisioning tools that let you create storage pools, volumes, and host security groups. These tools provide a wide range of functionality, such as the following:

- ["Managing Storage Pools"](#) on page 261
- ["Managing Volumes"](#) on page 263
- ["Rules for Creating Host Security Groups"](#) on page 270
- ["Managing Host Security Groups"](#) on page 275

About the Events Tab

The Events tab lets you view, clear, sort and filter events for an element. An event can be anything that occurs on the element, such as a device connected to a Brocade switch has gone off-line. It provides the following information about the events:

- **ID** - The identification number assigned to the event
- **Severity** - Provides the severity level
- **Time** - The time the event was recorded.
- **Summary Text** - A brief explanation of the event. When you click the summary text, the details of the event are displayed.

The Events tab lets you use Event Monitoring for Storage Essentials to:

- **View Event Details** - See "[Viewing Event Details](#)" on page 326.
- **Clear Events** - See "[Clearing Events](#)" on page 327.
- **Unclear Events** - See "[Unclearing Events](#)" on page 327.
- **Delete Events** - See "[Deleting Events](#)" on page 329.
- **Sort Events** - See "[Sorting Events](#)" on page 329.
- **Select a Severity for Filtering** - See "[About Filtering Events](#)" on page 333 and "[Selecting a Severity Level](#)" on page 333.
- **Customize the Severity Level Filter** - See "[Customizing the Severity Level Filter](#)" on page 334 and "[About Filtering Events](#)" on page 333.

Asset Attributes of an Element

IMPORTANT: Depending on your license, Chargeback Manager may not be available. See the "List of Features" to determine if you have access to Chargeback Manager. The "List of Features" is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).

Access asset information for an element by doing the following:

- Clicking an element (except for a file server) in Application Viewer, and then clicking the **Asset Management** tab.
- Double-clicking an element in Capacity Manager, Performance Manager or System Manager, and then clicking the **Asset Management** tab.
- Clicking an element in Chargeback.

Chargeback Manager provides a handy way for you to keep track of your asset information for an element. You can easily store warranty and licensing information, as well as contact information for the element. For example, assume a switch on the network is having some problems, and you want to contact the person in charge of that switch. You can use the element's asset record to not only find the contact information for that switch, but also the location of that switch.

The Asset Management tab displays general asset information about an element. It also provides access to other screens that provide additional asset information, such as staff, geographic, licensing, and warranty information. You can access these other screens by expanding the Asset Record node and clicking one of its children, as shown in the following figure. To learn more about these other screens, see the following topics:

- "[Adding Asset Information](#)" on page 463
- "[Adding General Information](#)" on page 464
- "[Adding Staff Information](#)" on page 465
- "[Adding Geographic Information](#)" on page 466
- "[Adding Licensing and Warranty Information](#)" on page 466

- ["Adding Custom Information"](#) on page 466

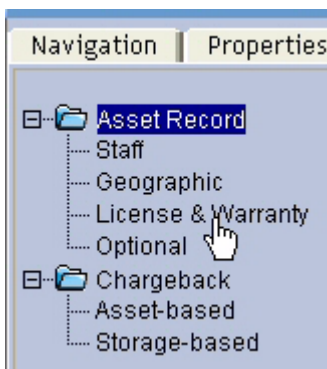


Figure 49 Viewing Asset Records

To set up chargeback, expand the Chargeback Manager node and click Asset-based or Storage-based. To learn more about each type of Chargeback Manager, see ["Setting Up Asset-Based Chargeback Manager"](#) on page 468 and ["Setting Up Storage-Based Chargeback Manager"](#) on page 471.

This section describes the fields on the Asset Management tab. When you are done with adding information on this page, click the **Save Changes** button at the bottom of the page.

NOTE: The fields that accept input cannot contain more than 250 characters.

- **Custom Name** - A name you assigned to the element. See ["Assigning a Custom Name"](#) on page 221 for more information.
- **Date Created** - Date the element was discovered.
- **Date Last Modified** - Date the record was last modified.
- **Description** - A description of the element. This description cannot be more than 250 characters.
- **Status** - The current status of the element. If the status of the element has changed, select the new status from the Status drop-down menu.
 - **New** - Default category for all detected elements.
 - **Missing** - No longer detectable through discovery
 - **Repaired** - The element is being repaired. The software does not automatically select this status.
 - **In Use** - The element is in use.
- **Vendor** - The vendor for the element.
- **Model** - The model of the element.
- **Serial Number** - Serial number of the element.
- **Barcode Number** - The barcode on the device.
- **Asset Code** - The asset code assigned to the element.

- **Asset Type** - The asset type assigned to the element.
- **Asset Tag** - The asset tag assigned to the element.
- **Asset Category** - The asset category assigned to the element.
- **Geographic Location** - The location of the element, for example, Boston, Massachusetts.
- **(Storage Systems Only) Storage Tier Classification** - Select the storage tier you want to assign the storage system. Storage tiers for storage-based information can have any name. By default, they are the following:
 - **No Tier** is the default setting. The storage is not classified.
 - **Ultra High Availability** is usually assigned to the ultra high availability storage. This tier contains the premium storage in your network, usually the most expensive.
 - **High Availability** is usually assigned to the high availability storage. This tier contains storage that is not as expensive as the storage assigned to the Ultra High Availability tier.
 - **IDE Based Storage** is usually assigned to IDE-based storage. This tier contains storage that is comparatively inexpensive.

About the Collectors Tab

The management server uses collectors to gather information. The **Collectors** tab provides information about the collectors for a particular element.

To start collectors and view reports for an element:

1. Access the **Collectors** page by doing one of the following:
 - Clicking an element in Application Viewer, and then clicking the **Collectors** tab.
 - Double-clicking an element in Capacity Manager, Performance Manager or System Manager, and then clicking the **Collectors** tab.
 - Clicking an element in Chargeback Manager, and then clicking the **Collectors** tab.
2. To change a collector's start time, modify the time and date entered in the **Next Scheduled Run** field. If you decide to change the start time, make sure the date is in the yyyy-mm-dd format with the time resembling a 24-hour clock. There should be a space between the date and the time, as shown below:


```
2005-06-26 09:41
```

After the collector runs, the value in this column is updated to the next time the collector will run.
3. To change how often the collector runs, type the number of minutes in the **Interval** field.

Important: Do not make the interval too short. Running a collector too frequently uses up space on the management server and impacts its performance.
4. To enable the collector, click the **Start** button.
5. To stop a collector, click the **Stop** button.
6. To view a report, click its link. See "[Viewing Reports](#)" on page 340 for more information.

About the Monitoring Tab

You can easily access performance information about an element by doing the following:

1. Double clicking the element in System Manager or Application Viewer.

2. Click the Monitoring tab.
The element appears highlighted in Performance Manager.
3. Select one of the monitoring options in the lower pane to view specific performance data about the element.
4. See "[Viewing Performance Data](#)" on page 399 for more information about Performance Manager.

About the Policies Tab

The Policies tab lets you view the utilization policies for an element. Utilization policies can automatically send an e-mail, generate an event, or run a custom script when an element is being overused. If the policy table is unpopulated, no policies exist for the element.

The Policies tab lets you use Policy Manager to do the following. See "[About Policy Manager](#)" on page 437 for more information:

- **Add Policies**
- **Test Policies**
- **Edit Policies**
- **Delete Policies**

You can access the Policies tab by doing one of the following:

- Double-clicking an element in Capacity Manager, Performance Manager or System Manager and then clicking the **Policies** tab.
- Right-clicking an element in Capacity Manager, Performance Manager or System Manager and then selecting **Show Policies** from the drop-down menu.

To access utilization policies for other elements and to create other types of policies, click the **Policy Manager** button in the left pane.

Determining If a Host Belongs to a File System

You can determine if a host is a member of a file system, such as CXFS™ on the Navigation tab and/or in Capacity Manager.

To find the information on the Navigation page:

1. Access System Manager as described in "[Accessing System Manager](#)" on page 163.
2. Double-click the host you want to determine is a member of a file system.
3. Click the Navigation tab.
4. Click the **Volumes** button.
5. Look under the File System Type column. The system type, such as CXFS, is listed.

Additional information about the storage volume is provided:

- Name of the storage volume
- Description of a storage volume
- Drive Type

To find the information in Capacity Manager:

1. Access Capacity Manager as described in "[Accessing Capacity Manager](#)" on page 425.
2. Select the host you want to determine is a member of a file system.
3. Scroll to the bottom of the page.
4. If a storage volume is a member of a shared file system, such as CXFS or XFS, it is displayed in the Storage Volume column.

You may need to expand the Storage Volume column if the volume names are long.

About the Data from CXFS File Systems

The management server can only monitor CXFS file systems from the host generating the input/output. For example, assume the elements in the following figure are part of a CXFS file system. When you generate input/out into the metadata server into /folder, only the metadata server is able to monitor the file system. For example, assume the metadata server generates 100 KB write, the management server displays 0 KB write for /folder on the metadata client.

The information in /folder on the metadata server is actually being mirrored to the /folder on the metadata client. The management server, however, does not detect the changes being mirrored to /folder on the metadata client.




Figure 50 CXFS File System

9 Provisioning Manager

IMPORTANT: Depending on your license, Provisioning Manager may not be available. See the “List of Features” to determine if you have access to Provisioning Manager. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).

This chapter describes the following:

- “[About Provisioning Manager](#)” on page 239
- “[Managing Zones](#)” on page 240
- “[Managing Storage](#)” on page 256

About Provisioning Manager

The software provides the following tools to assist you in provisioning your storage. These tools are accessible by clicking **Tools > Storage Essentials > Provisioning Manager**.

- **Path Provisioning tool** - Lets you schedule provisioning tasks to take place when the network traffic is light. See “[About Path Provisioning](#)” on page 289 for more information.
- **SAN Zoning tool** - Lets you create and modify zones, zone aliases and zone sets. Click the **Provisioning** button next to the fabric on which you want to do provisioning.

You can also view the properties of a fabric or switch by clicking its link in the table. See “[SAN Zoning Overview](#)” on page 240 for more information.

- **Storage System Provisioning tool** - Lets you manage storage pools, volumes, and host security groups. Click the **Provisioning** button next to the storage system on which you want to do provisioning.

You can also view the properties of a storage system by clicking its link in the table. See “[Setting Up Storage Partitioning](#)” on page 256.

IMPORTANT: Ports designated as an Initiator on a storage system belonging to the HDS Freedom Storage™ Lightning 9900™ Series or Freedom Storage Lightning 9900V Series cannot be used for provisioning. If you select one of these ports, you receive a message saying that provisioning failed because the HiCommand Database was not refreshed. The management server does not support provisioning for ports designated as Initiators on these storage systems.

Once you have become adept at provisioning, you might want to try accessing the provisioning screens by doing one of the following:

- Double-clicking a storage system or switch in System Manager and then clicking the **Provisioning Manager** tab.
- Right-clicking a storage system or switch in the Access tab in System Manager. See “[About the Access Tab](#)” on page 169 for more information.

Managing Zones

This section describes the following:

- ["SAN Zoning Overview"](#) on page 240
- ["Accessing Information About Zone Aliases"](#) on page 244
- ["Creating Zone Aliases"](#) on page 245
- ["Modifying a Zone Alias"](#) on page 246
- ["Deleting Zone Aliases"](#) on page 246
- ["Accessing Information About Zoning"](#) on page 247
- ["Creating a Zone in a Fabric"](#) on page 247
- ["Adding and Removing Zone Members"](#) on page 248
- ["Deleting Zones"](#) on page 249
- ["Accessing Information About Zone Sets"](#) on page 249
- ["Creating a Zone Set"](#) on page 249
- ["Modifying a Zone Set"](#) on page 251
- ["Deleting Zone Sets"](#) on page 251
- ["Copying a Zone Set"](#) on page 252
- ["Activating a Zone Set"](#) on page 253
- ["Zones and Zone Sets Are Sometimes Listed Twice"](#) on page 254
- ["Changing the Amount of Information Collected from the Inactive Zone Database \(Cisco Switches\)"](#) on page 255

SAN Zoning Overview

IMPORTANT: Depending on your license, Provisioning Manager may not be available. See the "List of Features" to determine if you have access to Provisioning Manager. The "List of Features" is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).

Use SAN zoning to control what can be seen in the storage area network (SAN). SAN zoning lets you group elements into zones, which can then be grouped into active and inactive zone sets. Only elements in an active zone set can be seen. A switch fabric can have multiple zone sets, but only one zone set can be active.

Zones are an excellent way to split hardware resources because they work by exclusion. For example, you can set up your switch ports so that elements connected to some of the ports appear in one zone and the rest appear in another zone. Members of a zone can only communicate with other members in the zone. If two elements are not within the same zone, they cannot communicate.

Zones are usually created for a particular task, such as controlling access between devices or groups. You might create zones based on an application or an operating system. For example, some network administrators prefer to put all of the Microsoft Windows computers in one zone and all of the Sun Solaris computers in another. As mentioned previously, you can create zones

according to an application. For example, you might want to create a zone for production and another zone for finance. This way the users in the finance department are not even aware of the disks and ports available for production and vice versa.

As shown in the following figure, ports 1 through 5 on the switch are assigned to a zone for production and ports 4 through 8 are assigned to the zone for the finance department. Users in the finance department can access storage systems B and C but not storage system A. Likewise, users in the production department can access storage systems A and C, but not storage system B. The following figure is an example of switch port zoning because the zones are assigned to ports on a switch.

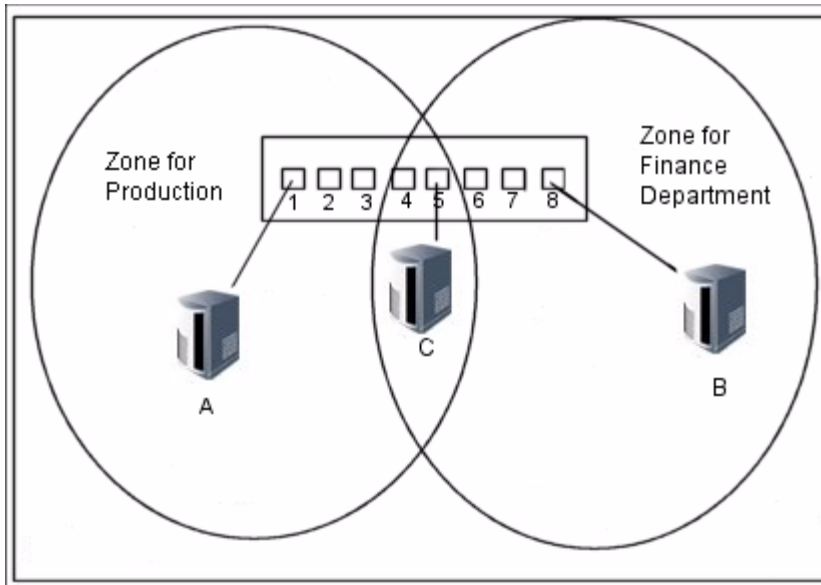


Figure 51 Resources in Two Zones

Only elements in an active zone set can communicate with each other. When a zone set is not active, it does not have any effect. If we do not want users in the Production and Finance zones to have access to the same storage, these two zones must be in two different zone sets, which must be both active. Since you can only have one active zone set to a fabric, the Production zone belongs to a zone set in one fabric and the Finance zone belongs to another zone set in another fabric.

A zone can be in more than one zone set. For example, the Finance zone could be in an active and inactive zone set. This allows for more flexibility. For example assume the Finance zone is a member of an active zone set named Zone Set One and a member of an inactive zone set named Zone Set Two in the same fabric. Zone Set Two contains additional zones. Assume you want to make users aware of those elements in the additional zones. When you activate Zone Set Two, users would still have access to the elements in Finance zone because it is also a member of Zone Two.

NOTE: Create zone aliases to easily keep track of your zones. Instead of remembering a port's name, you can assign a meaningful name. As a best practice, a zone should contain either zone aliases or ports, but not both.

The SAN Zoning tool is able to manage the two types of zoning:

- **Switch Port Zoning** - A switch port represented by a port number in a zone. Any device attached to the port is automatically in the zone.
- **WWN Zoning** - A device port represented by the Worldwide Name (WWN) in the zone.

The following diagram provides an overview of zoning:




Figure 52 Overview of Zoning Capabilities

Description

- **Zone Sets** - A zone set is a collection of zones. You can have only one zone set active at a time in a fabric; however, you can have a zone in more than one zone set. Zones sets are usually created for a particular task.
- **Zone** - A collection of zone aliases and ports.
- **Ports** - The Worldwide Name of the port to which an element is connected. A port in a zone can be of two types:
 - Switch Port Zoning
 - WWN Zoning

Use the following table as a guideline for setting up zoning.

Table 40 Setting Up Zoning

Step	Description	Where to Find
1	(Optional) Create a zone alias. Zone aliases are used to give meaningful names to switch ports, HBA ports, or storage system ports.	" Creating Zone Aliases " on page 245

Table 40 Setting Up Zoning (continued)

Step	Description	Where to Find
2	Create zones. Zoning is the primary tool to constrain groups of SAN members. A zone defines a logical SAN that contains limited element membership. The only elements visible to members of a zone are other members of that zone.	" Creating a Zone in a Fabric " on page 247
3	Create zone sets. A zone set contains multiple zones.	" Creating a Zone Set " on page 249
4	Activate a zone set. A switch fabric can have multiple zone sets defined, but only one zone set can be active.	" Activating a Zone Set " on page 253

Keep in mind the following:

- If you use another product to make zoning changes, such as add a zone, you must perform Discovery Data Collection for the management server to be made aware of these changes.
- The management server creates a zone by finding the port or Worldwide Name of discovered elements. The management server cannot create a zone by Fibre Channel addresses. An active zone will appear empty in the user interface of the management server if you use third-party software to create the zone by Fibre Channel addresses.

Use the following table for more information about what is supported for each type of switch.

Table 41 Zoning Support

Switch Type	Active Zones	Inactive Zones	Aliases	Zone/Zone Set Provisioning	Port Statistics
Brocade	Y	Y	Y	Y	Y
*McDATA SWAPI to EFCM	Y	Y	N	Y	Y
*McDATA SNMP through proxy	N	N	N	N	Y
*McDATA SNMP to switches	Y	N	N	N	Y
Cisco SNMP	Y	N	Y	N	Y
CNT SMI	Y	Y	N	N	Y

Table 41 Zoning Support

Switch Type	Active Zones	Inactive Zones	Aliases	Zone/Zone Set Provisioning	Port Statistics
QLogic SNMP	Y*	N	N	N	Y

*Also applies to EMC Connectrix switches.

**Also applies to Sun StorEdge switches.

***SANbox-2 only

Issues for Sun StorEdge and QLogic Switches:

- Active zoning is not reported for Sun StorEdge and QLogic switches that have the SNMPv1 Agent. Switches that do report active zoning, do not supply information about inactive zones to the management server.
- If a zone alias for a Sun StorEdge or QLogic switch is a member of an active zone, the zone alias is not shown in the user interface, but its members are displayed as belonging to the active zone.

Issues for Cisco switches:

- You can view zones, zone sets and zone aliases on a Cisco switch; however, you cannot use the management server to create, modify or remove them from a Cisco switch.
- The management server groups active zone sets in all Virtual SANs (VSANs) in a fabric into a zone set called "ACTIVE", and the "ACTIVE" zone set is shown associated with the physical fabric. The members of the "ACTIVE" zone set (zones, zone sets, zone aliases) have the name of the VSAN prefixed to their name. For example, an active zone named "ZONE1" from a VSAN named "VSAN1" is displayed as a zone on the physical fabric with name "VSAN1:ZONE1".

Issues for McDATA and Connectrix switches:

- Only one client at a time can provision on a McDATA or Connectrix fabric. Since each fabric has a separate lock, you can perform simultaneous provisioning on two different fabrics. For example, you could perform provisioning by using the user interface and the CLI at the same time on two different fabrics. Simultaneous provisioning on the same fabric is not supported.
- The management server does not support enabled default zones on McDATA or Connectrix switches. When a default zone is enabled on a McDATA or Connectrix switch, it is not listed as part of the active zone set.

Accessing Information About Zone Aliases



The software provides a listing of zone aliases in a fabric. You can view the properties of the zone alias and its port from this page.

To obtain information about zone aliases in a fabric:

1. Click **Tools > Storage Essentials > Provisioning Manager** in HP Systems Insight Manager.

2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric on which you want to do provisioning.
4. Click **Step 1 Zone Alias**.

This page lists the zone aliases and their ports under the following columns:

- **Name** - Click the name of the zone alias to view its properties.
- **Ports** - In some instance, you may be able to click the link of a port to view its properties.
- **Active** - A check mark appears in the Active column if the zone alias is included in an active zone set.
- **Edit** - Click the **Edit** button () to edit an alias. See ["Modifying a Zone Alias"](#) on page 246 for more information.
- **Delete** - Click the  button corresponding to the zone alias you want to delete. See ["Deleting Zone Aliases"](#) on page 246 for more information.

To create a zone alias, click the **New Zone Alias** button. See ["Creating Zone Aliases"](#) on page 245 for more information.

Creating Zone Aliases

Zone aliases are used to give meaningful names to switch ports, HBA ports, or storage system ports.

To avoid remembering a port's Worldwide Name (WWN), add a port to a zone alias as described in the following steps:

1. Click **Tools > Storage Essentials > Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to create a zone alias.
4. Click **Step 1 Zone Alias**.
5. Click the **New Zone Alias** button.
6. In the **Zone Alias Name** field, type a name for the zone alias.

Naming Conventions for Brocade Switches:

- The name must contain 1 to 64 characters.
- The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).
- The name is case sensitive. For example, "ZoneAlias1" and "zonealias1" are different zone aliases.
- You cannot create a zone alias with the same name as an existing zone, zone alias or zone set. For example, if you create a zone alias named "new", you cannot give a zone, zone alias, or zone set the same name.
- The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign (\$).


Naming Conventions for McDATA and Connectrix Switches:

- The name can have a maximum of 64 characters.
 - The first character of a zone alias name must be a letter (A-Z, AZ).
 - A zone alias name cannot contain spaces.
 - Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).
 - All names must be unique and may not differ by case. For example, myzonealias and MyZoneAlias are considered to be the same zone alias.
7. Add ports to the zone alias by selecting a port in the **Potential Ports** pane.
 8. Remove ports from the zone by selecting them in the **Ports in the Zone Alias** pane and clicking the **Remove From Zone** button.
 9. Click **OK**.

Modifying a Zone Alias

Zone aliases are used to give meaningful names to switch ports, HBA ports, or storage system ports.

To modify a zone alias:

1. Click **Tools > Storage Essentials > Provisioning Manager**.
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to modify the ports.
4. Click **Step 1 Zone Alias**.
5. Click the **Edit** button ()
6. Add ports to the zone alias by selecting a port in the **Potential Ports** pane.
7. Remove ports from the zone by selecting them in the **Ports in the Zone Alias** pane and clicking the **Remove From Zone** button.

NOTE: To select all of the ports, select the check box next to the Port heading.


8. Click **OK**.

Deleting Zone Aliases

You cannot delete a zone alias if it is the only member in a zone.

To delete a zone alias:

1. Click **Tools > Storage Essentials > Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to delete a zone alias.
4. Click **Step 1 Zone Alias**.



5. Click the  button corresponding to the zone alias you want to delete.
6. When you are asked if you want to delete the zone alias, click **OK**.

Accessing Information About Zoning

To access information about zones and to be able to manage them, do the following steps:

1. Click **Tools > Storage Essentials > Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to access information about zoning.
4. Click **Step 2 Zone**.

This page lists the zones, their aliases and ports under the following columns:

- **Name** - Click the name of the zone to view its properties.
- **Zone Aliases** - Click the name of the zone alias to view its properties.
- **Ports** - In some instance, you may be able to click the link of a port to view its properties.
- **Active** - A check mark appears in the Active column if the zone is included in an active zone set.
- **Edit** - Click the **Edit** button () to edit a zone. See ["Adding and Removing Zone Members"](#) on page 248 for more information.
- **Delete** - Click the  button corresponding to the zone you want to delete. See ["Deleting Zones"](#) on page 249 for more information.

To create a zone, click the **New Zone** button. See ["Creating a Zone in a Fabric"](#) on page 247 for more information.

Creating a Zone in a Fabric

To learn why zones are so important, see ["SAN Zoning Overview"](#) on page 240. A zone must have at least one member.

To create a zone:

1. Click **Tools > Storage Essentials > Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to create a zone.
4. Click **Step 2 Zone**.
5. Click the **New Zone** button.
6. In the **Zone Name** field, type a name for the zone.

Naming Conventions for Brocade Switches:

- The name must contain 1 to 64 characters.
- The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).

- The name is case sensitive. For example, "Zone1" and "zone1" are different zones.
- You cannot create a zone with the same name as an existing zone, zone alias or zone set. For example, if you create a zone named "new", you cannot give a zone, zone alias, or zone set the same name.
- The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign (\$).


Naming Conventions for McDATA and Connectrix Switches:

- The name can have a maximum of 64 characters.
 - The first character of a zone name must be a letter (A-Z, AZ).
 - A zone name cannot contain spaces.
 - Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).
 - All names must be unique and may not differ by case. For example, myzone and MyZone are considered to be the same zone.
7. Add members to the zone by selecting a member in the **Potential Members** pane. A zone member can be a port attached to a switch, a Worldwide Name (WWN) or a zone alias. Keep in mind the following:
 - As a best practice, a zone should contain zone aliases only, and there should be a zone alias for each port/WWN.
 - You cannot create a zone with an existing name.
 8. Remove members from the zone by selecting them in the Member Name pane and clicking the **Remove From Zone** button.
 9. Click **OK**.

Adding and Removing Zone Members

IMPORTANT: A zone must have at least one member.

To add and remove zone members:


1. Click **Tools > Storage Essentials > Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to modify a zone.
4. Click **Step 2 Zone**.
5. Click the **Edit** button (.
6. Add members to the zone by selecting a member in the **Potential Members** pane. A zone member can be a port attached to a switch, a Worldwide Name (WWN) or a zone alias.
7. Remove members from the zone by selecting them in the **Zone Member Names** pane and clicking the **Remove From Zone** button.
8. Click **OK**.

Deleting Zones

You cannot delete a zone if it is the only member in one of the zone sets or if it is a member of an active zone set. If you want to delete a zone in an active zone set, first move the zone to an inactive zone set. Then, delete it.

If you are using EFC Manager to delete zones, see ["Changes in EFC Manager Requiring "Discovery Data Collection""](#) on page 523.

To delete a zone:



1. Click **Tools** > **Storage Essentials** > **Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to delete a zone.
4. Click **Step 2 Zone**.
5. Click the  button corresponding to the zone you want to delete.
6. When you are asked if you want to delete the zone, click **OK**.

Accessing Information About Zone Sets

To access information about zone sets and to be able to manage them, do the following steps:

1. Click **Tools** > **Storage Essentials** > **Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to access information about a zone set.
4. Click **Step 3 Zone**.

This page lists information about zone sets under the following columns:

- **Name** - Click the name of the zone set to view its properties.
- **Zones** - Click the name of the zone to view its properties.
- **Active** - To make a zone set active, select its corresponding **Active** option. When you select a zone set, you make elements outside of the zone set inaccessible. See ["Activating a Zone Set"](#) on page 253 for more information.
- **Edit** - Click the **Edit** button () to edit a zone set. See ["Modifying a Zone Set"](#) on page 251 for more information.
- **Delete** - Click the  button corresponding to the zone set you want to delete. See ["Deleting Zone Sets"](#) on page 251 for more information.

To create a zone set, click the **New Zone Set** button. See ["Creating a Zone Set"](#) on page 249 for more information.

Creating a Zone Set

To learn why zone sets are so important, see ["SAN Zoning Overview"](#) on page 240.

To create a zone set in a fabric:

1. Click **Tools > Storage Essentials > Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to create a zone set.
4. Click **Step 3 Zone Set**.
5. Click the **New Zone Set** button.
6. In the **Zone Set Name** field, type a unique name for the new zone set.

Naming Conventions for Brocade Switches:

- The name must contain 1 to 64 characters.
- The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).
- The name is case sensitive. For example, "ZoneSet1" and "zoneset1" are different zone sets.
- You cannot create a zone set with the same name as an existing zone, zone alias or zone set. For example, if you create a zone set named "new", you cannot give a zone, zone alias, or zone set the same name.
- The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign (\$).

Naming Conventions for McDATA and Connectrix Switches:


- The name can have a maximum of 64 characters.
 - The first character of a zone set name must be a letter (A-Z, AZ).
 - A zone set name cannot contain spaces.
 - Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).
 - All names must be unique and may not differ by case. For example, myzoneset and MyZoneSet are both valid individually, but they are not considered to be unique.
7. If you want to make the zone set active, select the option **Make this the active zone set**. Keep in mind the following:
 - Only one zone set can be active at a time. When you make a zone set active, the previous active zone set becomes inactive.
 - You cannot create a zone set with the same name as the name of an existing zone set.
 8. Add zones to the zone set by selecting a zone in the **Zones Not in this Zone Set** and clicking the greater than sign (>).

NOTE: Zones can be in multiple zone sets. For example, zone1 could be in several zone sets.

9. Remove zones from the zone set by selecting a zone in the **Zones in this Zone Set** and clicking the less than sign (<).
10. Click **OK**.

Modifying a Zone Set

To modify a zone set:

1. Click **Tools** > **Storage Essentials** > **Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to modify a zone set.
4. Do one of the following:
 - Click **Step 3 Zone Set**.
 - Click **Step 4 Activate Zone Set**.
5. Click the **Edit** button ()
6. Add zones to the zone set by selecting a zone in the **Zones Not in this Zone Set** and clicking the greater than sign (>).

Note: Zones can be in multiple zone sets. For example, zone1 could be in several zone sets. Zone sets are usually created for a particular task.
7. Remove zones from the zone set by selecting a zone in the **Zones in this Zone Set** and clicking the less than sign (<).
8. Click **OK**.


Deleting Zone Sets

The software does not display all elements in a zone set, such as quick loop and fabric assist elements. When you delete a zone set, all elements, including quick loop and fabric assist, which are not viewable in the software, are deleted.

Only the zone set is deleted, not the zones contained in the zone set. For example, assume Zone A is contained in two zone sets: one named Zone_Set_One and another named Zone_Set_Two. If you delete Zone_Set_One, the zone has not been deleted so it is still in Zone_Set_Two.

If you are using EFC Manager to delete zone sets, see "[Changes in EFC Manager Requiring 'Discovery Data Collection'](#)" on page 523.

To delete a zone set:

1. Click **Tools** > **Storage Essentials** > **Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to delete a zone set.
4. Do one of the following:
 - Click **Step 3 Zone Set**.
 - Click **Step 4 Activate Zone Set**.
5. Click the  button corresponding to the zone set you want to delete.
6. When you are asked if you want to delete the zone set, click **OK**.

Copying a Zone Set

This feature copies a zone set and all of its members, such as zones and zone aliases. You can use this feature to copy inactive and active zone sets. The newly created zone set is inactive.


The management server stops the copying process of an active zone set if the management server finds one of the following:

- An inactive zone set with the same name as the name entered for the copy.
- An inactive zone with the same name as an active zone, but they do not have the same content.

Active zones in a zone set have corresponding inactive zones for redundancy. If you attempt to copy an active zone set containing a zone that does not have a corresponding inactive zone, the management server creates an inactive zone with the same name as the active zone. The inactive zone is used as a backup for the active zone.

IMPORTANT: This feature is not supported for Brocade switches and any switches that do not support provisioning. Aliases in the zone set are not copied over for McDATA switches.

To copy a zone set:

1. Click **Tools > Storage Essentials > Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the Provision button corresponding to the fabric in which you want to copy a zone set.
4. Click **Step 3 Zone Set**.
5. Click the corresponding  button for the zone set you want to copy.
6. Enter the name of the new zone set. If you are copying an active zone set, make sure you do not enter the name of a pre-existing inactive zone set. See the following restrictions: See the following restrictions:

Naming Conventions for McDATA and Connectrix Switches:

- The name can have a maximum of 64 characters.
 - The first character of a zone set name must be a letter (A-Z, AZ).
 - A zone set name cannot contain spaces.
 - Valid characters are a-a, AA, 0-9, ^, -, _, and \$.
 - All names must be unique and may not differ by case. For example, myzoneset and MyZoneSet are both valid individually, but they are not considered to be unique.
 - (Brocade Switches) You cannot create a zone set with the same name as an existing zone, zone alias or zone set. For example, if you create a zone set named "new", you cannot give a zone, zone alias, or zone set the same name.
7. Click **OK**.
- The zone is copied.

Activating a Zone Set

You can only have one zone set in a fabric active at a time; however, you could have a zone in more than one zone set.

To make a zone set active:

1. Click **Tools > Storage Essentials > Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **SAN Zoning** tab.
3. In the right pane, click the **Provision** button corresponding to the fabric in which you want to activate a zone set.
4. Click **Step 4 Activate Zone Set**.
5. Select the corresponding **Active** option, as shown in the following figure.
6. (McDATA and Connectrix switches) The management server lets you create a backup copy of the zone set you want to activate. To create a backup of the zone set that will become active, do the following:
 - a. Select the option, **Make a backup copy of the active zone set after activation**.
 - b. (Optional) In the **Name** field, modify the name that has been assigned to the backup zone set. The management server assigns the name by appending the date and time of the zone set you have selected to become active, as shown in the following example:

`zone_name_2005-05-17_13-41-05`

where

- `zone_name` is the name of the zone you are making active.
- `2005-05-17` is the date you made the zone active. The date is formatted as year-month-day. In the example, the date is May 17, 2005.
- `13-41-05` is the time the copy was made. The time is formatted as hour-minute-second, and it uses the 24-hour notation. In the example, the time is 1:41:05 p.m.

The name of the backup zone set must follow the following naming conventions for McDATA and Connectrix switches:

- The name can have a maximum of 64 characters.
- The first character of a zone set name must be a letter (A-Z, AZ).
- A zone set name cannot contain spaces.
- Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).
- All names must be unique and may not differ by case. For example, myzoneset and MyZoneSet are both valid individually, but they are not considered to be unique.

NOTE: The management server truncates the name of the backup zone set if it is more than 44 characters long. For example, assume you have a zone set named `McDATA_Switches_Burlington_Massachusetts_United_States`. The management server truncates the name to `McDATA_Switches_Burlington_Massachusetts_Un_2005-05-17_13-41-05` when it creates the backup zone set because the original name was more than 44 characters long. The management server truncates the original name so that it can fit the date and time into the zone set name, which cannot be more than 64 characters.

7. To activate the zone, click **OK**.

IMPORTANT: Only one zone set can be active at a time. When you make a zone set active, the previous active zone set becomes inactive.

Zones and Zone Sets Are Sometimes Listed Twice

Sometimes the Navigation tab and the Provisioning Manager pages list the zones and zone sets twice for McDATA and Connectrix switches. EFC Manager and Connectrix Manager contain an off line Zoning Library. This Zoning Library holds all zone sets, zones and zone members. When you activate a zone set, the zone set along with its zones and zone members is copied to the McDATA and/or Connectrix switches and activated. This creates an active copy, in addition to the saved copy that already exists in the Zoning Library. If you edit the saved copy of the zone or zone set in the Zoning Library, you have to reactivate the saved copy of the active zone set.

For example, assume you have the following information in the EFC Zoning Library:

```
ZoneSet A
    Zone A1
        ZoneMember A1a

ZoneSet B
    Zone B1
        ZoneMember B1b
```

If a user activates ZoneSetB, the existing information in ZoneSetB is copied to the switch and activated. The Zoning Library, however, still contains the older information. The following is what you would see in the Zoning Library:

```
ZoneSets:
  ZoneSet A
  ZoneSet B    (this is the inactive ZoneSet in the Zoning Library)
  ZoneSet B    (this is the active ZoneSet which is not the same as the one
above)
Zones:
  Zone A1
  Zone B1    (this is the inactive one)
  Zone B1    (this is the active one)
ZoneMembers:
  ZoneMember A1a
  ZoneMember B1b    (inactive)
  ZoneMember B1b    (active)
```

Notice ZoneSetB and its members are listed twice in the previous example. They are also displayed twice on the **Navigation** tabs and Provisioning Manager pages. If you click Zone A1 in EFC Manager or Connectrix Manager, you are shown it has member ZoneMember A1a. The other zone, Zone B1, has ZoneMember B1b. The name of an item in the provider cannot be changed, because the management server is required to report the actual name of the element.

If you remove a zone from the Zoning library, the zone is shown only once in the user interface of the management server. The zone is not displayed at all in the management server if it is an inactive zone. Each time you make changes in the Zoning Library using EFC Manager or Connectrix Manager, you must run Discovery Data Collection (**Discovery > Details**) for the management server to obtain the latest information.

Changing the Amount of Information Collected from the Inactive Zone Database (Cisco Switches)

You can change how much information the management server collects during Discovery Data Collection from the inactive database for Cisco switches. The zone database is stored on each switch. If the switches are configured to distribute zone database changes to all switches in the fabric, it is only necessary to get zone database information from one switch in each VSAN.

The following options are supported:

- **all** - Show all zone database information from all switches. This may produce duplicate information and result in large zone tables if zone databases are distributed and/or large.
- **primary (default)** - Show only zone database information from the primary switch in each VSAN. This option is the recommended setting if zone information is distributed to all switches in the fabric or if only the primary switch zone database is used. Note that the primary switch can be different for different VSANs.
- **none** - Do not display any zone database information. This option is not recommended.

To change the amount of information collected:

1. Click **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.

```
cimom.cisco.displayZoneDatabase=primary
```

4. Return to the Advanced page (**Options > Storage Essentials > Manage Product Health**). Then, click **Advanced** in the Disk Space tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Change the value of `cimom.cisco.displayZoneDatabase`. Make your changes in the **Custom Properties** field. For example, to set the value of `cimom.cisco.displayZoneDatabase` to `all`, you would change the value from `primary` to `all` as shown in the following example:

```
cimom.cisco.displayZoneDatabase=all
```
7. When you are done, click **Save**.
8. Restart the service for the management server for your changes to take effect.

While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Managing Storage

This section describes the following:

- ["Setting Up Storage Partitioning"](#) on page 256
- ["Modifying the Cache Settings \(Engenio Only\)"](#) on page 260
- ["Changing the Owner of a Volume \(Engenio Only\)"](#) on page 260
- ["Managing Storage Pools"](#) on page 261
- ["Managing Volumes"](#) on page 263
- ["Rules for Creating Host Security Groups"](#) on page 270
- ["Managing Host Security Groups"](#) on page 275
- ["General Provisioning Issues"](#) on page 282
- ["Provisioning Issues by Vendor"](#) on page 282

Setting Up Storage Partitioning

Each storage vendor treats storage partitioning differently. For example, Hitachi and EMC ship their storage systems with the volumes already created. Other storage vendors, such as Engenio ship their storage system as an empty array.

Despite the differences among storage systems, you can still use this product to manage your provisioning. Some tasks, such as volume creation, might create different results according to the

type of storage system. To learn how host security groups are created on your storage systems, see ["Rules for Creating Host Security Groups"](#) on page 270.

To learn which storage systems are supported, see [Table 42](#) on page 257 and [Table 43](#) on page 258.

Table 42 Provisioning and Pool Support

Storage System	Storage Provisioning	Create/Delete Pool	Create Pool Using Settings ³
Engenio and Sun 6130	Y	Y	Y
CLARiiON	Y, 1	Y	Y
Symmetrix	Y	N	N
HDS	Y	N	N
Sun 35xx	N	N	N
Sun 6920	Y	N	N
HP-XP SMI	N	N	N
HP-MSA	N	N	N
HP-EVA	N	N	N
IBM ESS2	Y	N	N

Note from table:

1The EMC Navisphere® CLI is required to communicate with the CLARiiON storage system. The CLARiiON storage system must be configured to recognize the management server as a privileged user. See ["EMC Navisphere CLI Is Required"](#) on page 283 for more information.

2IBM and other storage systems that use external providers show raw capacity for a storage pool, instead of formatted capacity.

3The "Create Pool Using Settings" column refers to the functionality that lets you choose the type of pool, usually RAID level.

See the following table for information regarding support for volumes and host security groups.

Table 43 Volume and HSG Support

Storage System	Create/Delete Volume	Create Volume Using Settings, ⁶	Create/Delete Meta Volume	HSG Provisioning Supported
Engenio and Sun 6130	Y	Y	N	Y
CLARiiON	Y	Y, 5	Y	Y
Symmetrix	Y, 1	N	Y	Y, 2
HDS	Y, 1	N	Y	Y
Sun 35xx	N	N	Y, 3	N
7Sun 6920	Y	N	N	Y
HP-XP SMI	N	N	N	N
HP-MSA	Y	N	N	Y
HP-EVA	Y	N	N	Y
IBM ESS	Y, 4	N	N	Y

Notes from table:

1. Creating volume means marking existing device as accessible to HSG management, deleting volume means returning device to free device pool.
2. Solutions Enabler 5.0 or later is required.
3. LUSE volumes are not created until the management server meta volume is mapped in an HSG.
4. Volumes created on IBM storage systems cannot be deleted.
5. RAID level can be specified for first volume in a pool, subsequent volumes inherit this setting
6. The "Create Volume Using Settings" column refers to the provisioning capability on the management server that lets you create different kinds of volumes from a pool depending on whether you want it optimized, such as for streaming, random access or high availability.
7. Unlike the native tool for the Sun 6130 storage system, the management server refers to virtual disks as storage pools.

See the following topics for more information about how storage provisioning works on your storage system.

- ["Issues Specific to CLARiiON Storage Systems"](#) on page 283
- ["Issues Specific to EMC Symmetrix Storage Systems"](#) on page 283
- ["Issues Specific to HDS Storage Systems"](#) on page 285
- ["Issues Specific to Engenio Storage Systems"](#) on page 287

The following table provides an overview of the required steps for setting up storage partitioning.

Table 44 Setting Up Storage Partitioning

Step	Description	Where to Find
1	<p>(Engenio and CLARiiON® only) Create a storage pool. Sometimes referred to as a volume group or RAID group.</p> <p>A storage pool is a group of disks associated together through a RAID configuration. The pool's capabilities define the level of protection for the associated volumes and LUNs.</p> <p>EMC Symmetrix and HDS storage systems have storage pools that are predefined. Engenio and CLARiiON storage systems require a storage pool to be created and volumes to be allocated from a storage pool.</p>	"Creating a Storage Pool (Engenio and CLARiiON Only)" on page 261
2	<p>Create a volume.</p> <p>A volume is a virtual disk. Volumes are created in sizes that are desirable for being shown as a LUN.</p> <p>On CLARiiON storage systems, a volume is owned by one of the storage processors. Creating a volume also creates a LUN for this volume and each port of the storage processor that owns this volume. A volume can be associated with more than one fibre channel port, creating multiple LUNs corresponding to the same volume. The defining characteristics of a LUN are the volume, port, and LUN number. On CLARiiON storage systems, a LUN mapped to a port is visible to all the ports on that controller. Mapping a volume to a port on a CLARiiON storage system also maps that volume to all ports that reside on the same storage processor as the selected port. It also causes the volume to be unmapped from all the ports of the other storage processor.</p> <p>Some storage systems have their volumes fully configured during install. For these storage systems, Users can concatenate multiple volumes together to create a new volume.</p>	"Creating a Storage Volume" on page 265


Table 44 Setting Up Storage Partitioning (continued)

Step	Description	Where to Find
3	Create a Host Security Group. Host Security Groups define which initiators (HBA ports) have access to specified storage volumes. They are associated with a fibre-channel port and contain a list of HBA port initiators and the volumes they can see.	" Creating Host Security Groups " on page 276 and " Rules for Creating Host Security Groups " on page 270

Modifying the Cache Settings (Engenio Only)

IMPORTANT: Depending on your license, Provisioning Manager may not be available. See the "List of Features" to determine if you have access to Provisioning Manager. The "List of Features" is accessible from the Documentation Center (**Help** > **Documentation Center** in Storage Essentials).

To modify the cache settings:

1. Click the  button corresponding to the volume you want to modify.
2. Type the cache read ahead multiplier (0 to 65535 bytes) in the **Cache read-ahead multiplier** field.
A cache read ahead multiplier copies additional data blocks into the cache while it is reading and copying host-requested data blocks from disk to cache. To disable this option, type 0.
3. (Optional) Select **Read Caching**.
When this option is enabled, the host's operations are stored in controller cache memory.
4. (Optional) Select **Write Caching**.
When this option is selected, data is written to the cache memory of a controller
5. (Optional) Select **Write Caching with Mirroring**.
Use this option to preserve data if a controller or the cache fails. When this option is enabled, the data is written to two redundant controllers of the same cache size. This configuration provides redundancy in case a controller fails. One controller performs uncompleted write operations when the other controller fails.


NOTE: See "[Changing the Owner of a Volume \(Engenio Only\)](#)" on page 260 for information about changing the owner assigned to the volume.

Changing the Owner of a Volume (Engenio Only)

When a volume is created, the management server automatically assigns a controller to be the owner of the volume. You can change the owning controller if you want to use a different one for LUN masking.

NOTE: If the owner becomes unreachable as a result of a network failure or the owner itself fails, the other controller in the pair automatically becomes the owner of the volume.

To change the owner:

1. Click the  button corresponding to the volume you want to modify.
2. Select the owner for the volume from the **Current Owner** drop-down menu.
3. Click **OK**.

Managing Storage Pools

This section describes the following:

- “[Creating a Storage Pool \(Engenio and CLARiiON Only\)](#)” on page 261
- “[Accessing Information About Storage Pools](#)” on page 262
- “[Deleting a Storage Pool \(Engenio and CLARiiON Only\)](#)” on page 263

Creating a Storage Pool (Engenio and CLARiiON Only)

A storage pool is a group of disks associated together through a RAID configuration. The pool’s capabilities define the level of protection for the associated volumes and LUNs. Create at least one storage pool before provisioning a volume.

To create a storage pool:

1. Click **Tools** > **Storage Essentials** > **Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **Storage Systems** tab.
3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to create a storage pool.
4. Click **Step 1 Storage Pool**.
5. Click the **New Storage Pool** button.
6. Select a setting and size for the storage pool.
The sizes displayed depends on the RAID level you want. For example, RAID 0 does not require additional drives so you can assign more space to the pool.
7. Click **OK**.
The storage pool is created.

NOTE: When you create a pool on an Engenio storage system, a placeholder volume is created inside the new volume group. The name of the placeholder volume starts with “Required - do not delete.” The placeholder volume is required because the storage pool cannot not exist without it. The management server does not display the placeholder volume, but other monitoring products may display this volume.


Accessing Information About Storage Pools

If you use another product to make provisioning changes, you must perform Discovery Data Collection (**Discovery** > **Details**) for the management server to be made aware of these changes.

To access information about storage pools:

1. Click **Tools** > **Storage Essentials** > **Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **Storage Systems** tab.
3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to access information about storage pools.
4. Click **Step 1 Storage Pool**.

This page lists information about the storage pools under the following columns:

- **Pool Name** - Click the name of the storage pool to view its properties.
- **Size** - Displays the amount of space assigned to the storage pool.
- **Available** - Displays the amount of space available in the storage pool.
- **Used** - Displays the amount of space used in the storage pool.
- **Volumes** - Click the name of the volume to view its properties.
- **Capabilities** - Click the RAID level or name to view its properties.
- ***Mainframe** - Displays whether the storage pools have volumes that are on a mainframe.
- ****Delete** - Click the  button corresponding to the storage pool you want to delete. See the topic, “[Deleting a Storage Pool \(Engenio and CLARiiON Only\)](#)” on page 263 for more information.

NOTE: To create a storage pool, click the **New Storage Pool** button in the upper-right corner of the page. See the topic, “[Creating a Storage Pool \(Engenio and CLARiiON Only\)](#)” on page 261 for more information.


*This option is available to only HDS storage systems.

**These options are available to only Engenio and CLARiiON storage systems.

Deleting a Storage Pool (Engenio and CLARiiON Only)

IMPORTANT: When you delete a storage pool on an Engenio storage system, all the volumes for the volume group are deleted, including the placeholder volume.

To delete a storage pool:

1. Click **Tools** > **Storage Essentials** > **Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **Storage Systems** tab.
3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to delete a storage pool.
4. Click **Step 1 Storage Pool**.
5. Click the  button corresponding to the storage pool you want to delete.
When you delete a storage pool, you are warned about the other volumes that will be deleted.
6. Click **OK**.
7. The storage pool and its volumes are deleted.

Managing Volumes

This section describes the following:

- ["Accessing Information About Volumes"](#) on page 263
- ["Creating a Storage Volume"](#) on page 265
- ["Deleting a Storage Volume"](#) on page 268
- ["Changing the Cache Block Size for a Storage System"](#) on page 269
- ["Modifying the Cache Settings \(Engenio Storage Systems Only\)"](#) on page 269

Accessing Information About Volumes

IMPORTANT: Some storage vendors require a password to access the storage system. If the same password is not entered, an authentication error message is displayed. Refer to **Discovery** > **Details** to determine the user account that was used to access the storage system during discovery.



If you use another product to make provisioning changes, you must perform Discovery Data Collection (**Discovery** > **Details**) for the management server to be made aware of these changes.

To access information about volumes:

1. Click **Tools** > **Storage Essentials** > **Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **Storage Systems** tab.
3. Click **Step 2 Volume**.
The right pane appears blank.
4. To view volumes, do the following:

- a. Select one of the following:
 - **All Volumes** - Displays all volumes.
 - **Unmapped Volumes** - Displays all unmapped volumes.
 - **Mapped Volumes** - Displays all mapped volumes.
 - **A port under the Ports folder** - Displays the volumes that use the selected port.
 - **A storage pool under the Storage Pools folder** - Displays the volumes that are part of the storage pool.
- b. (Optional) To narrow your search, you can specify the following:
 - **A string** - To display volumes containing a string, enter the string in the **Name Contains** field.
 - **Size** - To display volumes with the specified size in megabytes, enter the size in the **Size** field.
- c. Click **Apply**.

The following columns list information about the volumes:

- **Volume** - Click the name of the volume to view its properties.
- **Size** - Displays the size of the volume in megabytes (MB).
- **Ports** - Click a ports link to view its properties.
- **Pool** - Click the name of the pool to view its properties.
- ***Mainframe** - Displays whether the storage pools have volumes that are on a mainframe.
- ****Default Owner** - The controller which owns the storage system when it is rebooted.
- ****Current Owner** - The controller which currently owns the system.
- ****Segment Size** - Displays the amount of space assigned to a volume in megabytes (MB).
- ****Read ahead** - Displays the cache read ahead multiplier.
- ****Edit** - Click the  button corresponding to the volume you want to edit. See the topic, "Modifying the Cache Settings," for more information.
- **Delete** - Click the  button corresponding to the volume you want to delete. See the topic, "Deleting a Storage Volume" on page 268 for more information.

*This option is available to only HDS storage systems.

**Not accessible to all storage systems.

To create a volume, click the **New Volume** button in the upper-right corner of the page. To delete several volumes at once, select the volumes you want to delete and then click the **Delete Selected Volumes** button.

See "[Volumes on HDS Storage Systems](#)" on page 264 if you have an HDS storage system.

Volumes on HDS Storage Systems

Volumes from single LDEVs are shown as **LDEV: 0** on HDS storage systems. When volumes made up of multiple LDEVs are first created, they are not mapped to a target port on the storage system. The software remembers these LDEVs constitute a single volume, but it does not make changes to

the storage system until the volume is mapped to a port. As a result, they are referred to as Groups, for example:

```
Group:0 (LDEV:0, LDEV:1)
```

where 0 in Group:0 is the volume identifier and LDEV:0 and LDEV:1 are the LDEVs that make up this volume.

After you create a storage volume on an HDS storage system, you must map the volume to a target port on the storage system. You can map the volume to the target port by using the storage system Provisioning tool. In the tool, click **Step 3, LUNs**.

Once the volume is mapped, it is displayed as a logical unit size expansion (LUSE), as shown below:

```
LUSE:0 (LDEV:0, LDEV:1)
```

where 0 in Group:0 is the volume identifier and LDEV:0 and LDEV:1 are the LDEVs that make up this volume.

Creating a Storage Volume

IMPORTANT: Some storage vendors require a password to access the storage system. If the same password is not entered, an authentication error message is displayed. Refer to **Discovery > Details** to determine the user account that was used to access the storage system during discovery.

When you create a storage volume, you can set its size, volume capacity, and storage pool.

Volumes are shipped already created on HDS and Symmetrix storage systems. When you create a volume in the management server on these storage systems, you are defining the volume as being allocated.

Keep in mind the following:

- Some vendor's tools for HDS might round off the volume size, so that a 6.87-GB volume appears as 7 GB (7168 MB) in the tool. The management server displays the size of the volume without rounding. For example, assume you want to create a 14-GB (14336 MB) LUSE volume, and according to the storage tool, you have two 7-GB LDEVs, which are really 6.87 GB (7034.88 MB). If you look at the native tool, it would be logical to assume only two LDEVs would be used to create the 14-GB LUSE volume. The management server would use three LDEVs because each LDEV is really 6.87 GB.
- If you are creating a volume on an HP EVA storage system, its external SMI-S provider may round the specified number of megabytes to the nearest whole gigabyte.
- You cannot create volumes in existing storage pools on HP MSA and EVA storage systems.
- PvLinks based on HP disk partitions are not supported. Any volumes created on such PvLink meta devices are shown as local. If you partition a regular (non-PvLink) external disk and create volumes based on it, then volumes are recognized as external volumes.

Keep in mind the following for HDS storage systems:

- A LUSE volume on an HDS storage system is not created until you map that volume to a target array port. In the management server, the create volume and LUN creation tasks are two different operations. So if you want to create a LUSE volume and perform LUN creation on the HDS box, it is a two-step process. First, use the management server to create LUSE volumes. The management server puts volumes in the repository as grouped LDEVs. Then, create a LUN and map a volume to a target port, by creating a host security group. See the topic, "[Creating Host Security Groups](#)" on page 276 and "[Rules for Creating Host Security Groups](#)" on page 270.
- HDS ships some of its storage systems with volumes already created. When the software first discovers an HDS storage system, it detects the volumes created by HDS. When you use the software's "create a volume" feature, you are assigning the already created volume. See the topic, "[About Provisioning on HDS Storage Systems](#)" on page 285 for more information.
- LUSE made up of volumes from different RAID levels are not supported. You cannot use this product to provision this type of LUSE. Existing LUSE of this type maybe incorrectly reported.

Keep in mind the following for Engenio storage systems:

- Make sure you select a volume group that can accommodate the requested size for the new volume.
- You can create volumes from existing free extent areas within a volume group.
- No volume-to-LUN masking is done by default, except for CLARiiON storage systems. See the topic, "[Creating Host Security Groups](#)" on page 276 and "[Rules for Creating Host Security Groups](#)" on page 270.

To create a storage volume:

1. Access the Create Storage Volume wizard by doing the following:
 - a. Click **Tools > Storage Essentials > Provisioning Manager** in HP Systems Insight Manager.
 - b. In the right pane, click the **Storage Systems** tab.
 - c. In the right pane, click the **Provision** button corresponding to the storage system in which you want to access information about volumes.
 - d. Click **Step 2 Volume**.
 - e. If you want to create LUSE volumes on an HDS storage system, click the **Unmapped Volumes** folder in the left pane. Then, select the desired number of LDEVs for the LUSE volume then click the **Delete Selected Volumes** button. Take note of the array group from which you deleted the LDEVs. You need this information to create the LUSE volume.
2. Click the **New Volume** button.

NOTE: You can also access the Create Storage Volume wizard from the Navigation tab in System Manager. To access the wizard from the Navigation tab, click the Volumes link for a storage system, then click the **New Volume** button at the bottom of the screen.

3. (Engenio only) In the **Volume Name** field, type a name for the volume. If you do not provide a name, the software assigns one.

If you type a volume name, keep in mind the following:

- It cannot be more than 30 characters.
 - The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or one of the following symbols: dollar sign (\$), caret (^), or an underscore (_).
 - The name is case sensitive, for example, "StorageVolume1" and "storagevolume1" are different storage volumes.
 - The name must be different from any other volume name on an Engenio storage system.
4. In the **Size** field, type the size of the volume in megabytes (MB) or gigabytes (GB). Select the appropriate unit of measurement from the drop-down menu to the right of the **Size** field.

IMPORTANT: The management server creates a volume of at least the size specified. For example, assume you requested that the management server create a 15-MB volume, and you have only three free extents: 10 MB, 17 MB, and 100 MB. A 17-MB volume will be created instead of a 15-MB volume because that is the closest size you have for a free extent. Keep in mind that although the management server tries to find free extents that make the volume size as close to the requested size, it does not guarantee it will pick the optimal combination of free extents.

5. Select a storage pool for the volume.

Keep in mind the following:

- If you do not see space available in the storage pool, you must delete volumes. For example, assume you want to create an 8-MB volume, but you do not have space available. Each volume is made up of 4 MB. You must delete two volumes from that storage pool. Make sure those volumes you delete are not being used. See "[Deleting a Storage Volume](#)" on page 268 for more information.
- If you do not see a storage pool, verify that you have obtained all element details from the storage system. See "[Updating the Database with Element Changes](#)" on page 45 for more information.

6. Click **Next**.

7. Select a volume capability. The volume capabilities listed depend on the type of storage system. For example, if an EMC Symmetrix storage system is selected, **Pool default settings** are displayed in the field. If an Engenio storage system is selected, the following is displayed.

- **<Default>** - Provides the default cache read ahead multiplier and the default segment size for the storage system.
- **File System (Typical)** - Provides a cache read ahead multiplier of 1 with a segment size of 64 KB.
- **Database** - Provides a cache read ahead multiplier of 0 with a segment size of 64 KB.
- **Multimedia** - Provides a cache read ahead multiplier of 8 with a segment size of 128 KB.
- **Custom** - Lets you customize the cache read ahead multiplier and the segment size.

NOTE: (HDS only) Under the “Volume Capabilities Tab”, leave the default selection “<default>” selected and click the “Finish” button. Once the settings have been made, you see a new volume called a Group Volume in the list of unmapped volumes. Technically this is not a LUSE yet as it has not been assigned to a port; it is a logical grouping within the management server. Think of it as a place holder. From this point you can select the new group volume and assign it to a port. Once the volume has been assigned to a port, the management server makes the changes to the array and creates the LUSE. See [Table](#) on page 276 and [Table](#) on page 270 for information on how to create a LUN.

For information about the volume capability, refer to the documentation accompanying the storage system.

8. (Engenio only) If you selected the **Custom** option, do the following:

a. Type the cache read ahead multiplier (0 to 65535 bytes) in the **Cache read ahead multiplier** field.

A cache read ahead multiplier copies additional data blocks into the cache while it is reading and copying host-requested data blocks from disk to cache. Select the multiplier that maximizes performance for the way the volume will be utilized.

b. Select a segment size from the drop-down menu.

9. Click **Finish**.


Deleting a Storage Volume

When you delete a storage volume on an HDS or Symmetrix storage system, the software marks the deleted volume as hidden in the CIM repository, thus, making it unassigned, instead of being deleted. The software keeps track of the “deleted volumes.”

Keep in mind the following:

- Some storage vendors require a password to access the storage system. If the same password is not entered, an authentication error message is displayed. Refer to **Discovery > Details** to determine the user account that was used to access the storage system during discovery.
- If you remove volumes from host storage groups that are command devices or are pair volumes on HDS storage systems, later modification of the pair volumes may be disabled.

To delete a storage volume:

1. Click **Tools > Storage Essentials > Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **Storage Systems** tab.
3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to access information about volumes.
4. Click **Step 2 Volume**.
5. To display only unmapped volumes, click the unmapped volumes node in the left pane.
6. Click the  button corresponding to the volume you want to delete.
7. When you are asked if you want to delete the volume, click **OK**.

8. To delete several storage volumes at once, select the storage volumes you want to delete and then click the **Delete Selected Volumes** button.

NOTE: To select all volumes, select the check box next to the Volume heading, as shown in the following figure:

Changing the Cache Block Size for a Storage System


Engenio lets you change the cache block size on their storage systems.

To change the cache block size of a storage system:

1. Click **Tools > Storage Essentials > Provisioning Manager** in HP Systems Insight Manager.
2. Double-click the storage system displayed in the right pane.
3. Scroll to the bottom of the Navigation page.
4. Click the **Change** button.
5. Select the cache block size from the drop-down menu. Set a higher cache size for applications that requires a lot of input and output, such as multimedia.
6. Click **OK**.

Modifying the Cache Settings (Engenio Storage Systems Only)

To modify the cache settings:

1. Click **Tools > Storage Essentials > Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **Storage Systems** tab.
3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to access information about volumes.
4. Click **Step 2 Volume**.
5. Click the **Edit** button () corresponding to the volume you want to modify.
6. Type the cache read ahead multiplier (0 to 65535 bytes) in the **Cache read-ahead multiplier** field.

A cache read ahead multiplier copies additional data blocks into the cache while it is reading and copying host-requested data blocks from disk to cache. To disable this option, type 0.

7. (Optional) Select **Read Caching**.

When this option is enabled, the host's operations are stored in controller cache memory.

8. (Optional) Select **Write Caching**.

When this option is selected, data is written to the cache memory of a controller

9. (Optional) Select **Write Caching with Mirroring**.

Use this option to preserve data if a controller or the cache fails. When this option is enabled, the data is written to two redundant controllers of the same cache size. This configuration provides redundancy in case a controller fails. One controller performs uncompleted write operations when the other controller fails.

Rules for Creating Host Security Groups

The management server now uses host security groups instead of LUN masking and LUN mapping. With the introduction of host security groups, the management server has a new definition of mapped for this release. Mapped is capacity that is accessible by one or more hosts external to the array (aggregated capacity of volumes that are accessible from hosts external to the subsystem).

Each storage system treats host security groups differently, as shown in the following table.

Table 45 HSG Port Assignments

Storage System	Port Assignment
EMC CLARiiON	<ul style="list-style-type: none">• Volumes can be only on SPA or SPB because CLARiiON is active/passive storage, which means it can have only one active path to a volume. Addition of initiators to any of the processors is listed for both processors.• The HSG is created on all ports of the processor you select unless you select an initiator that uses a different processor and the initiator does not belong to a host security group. For example, assume you select processor SPA. Then, you select an initiator that belongs to SPB but it does not belong to a host security group. The host security group is created for all ports on SPB.• When you select an initiator for the host security group, the initiator has to be registered with the CLARiiON storage system.• You can have more than one initiator in a security group if you have the proper multipathing software installed on that particular host where the initiator is located.
Engenio	<ul style="list-style-type: none">• When you create the HSG by using the management server, you cannot specify the controller, but not the port for the HSG.• An initiator is equivalent to a host port in SANtricity.• You can have multiple volumes and initiators in the HSG.
HDS and HP XP storage systems using Command View XP Advanced Edition	<ul style="list-style-type: none">• The FC port contains only volumes but no initiators (HBA WWN) assignment, the management server displays these volumes as unmapped since no external host can see these volumes yet.• Every HSG is on one port on the array. You can have host security groups with the same name, as long as they are on different ports.

Table 45 HSG Port Assignments (continued)

Storage System	Port Assignment
IBM	<ul style="list-style-type: none"> • A volume can be assigned to more than one initiator, but you can have only one initiator in a host security group. • You can select any number of ports from one to all when creating the HSG. • You can create the HSG with or without LUNs. • The default is all ports are in the host security group. If no ports are selected, the default is used.
Sun 6920	<ul style="list-style-type: none"> • You can have two or more initiators assigned to the same volume, but you can have only one initiator to a host security group. • You must create the HSG with one initiator and at least one LUN. You can have two or more initiators assigned to the same volume. • You can select any number of ports from one to all when creating the HSG. • Creation of a new HSG with an initiator that already exists in another HSG adds the volumes to the previously created HSG. • The management server displays the ports connected to the initiator you selected for creating an HSG. If the initiator is not connected to any ports, no ports are displayed.
Sun 6130	<ul style="list-style-type: none"> • When you create the HSG by using the management server, you cannot specify the controller, but not the port for the HSG. • An initiator is equivalent to a host port in the native tool. • You can have multiple volumes and initiators in the HSG.
Symmetrix	<ul style="list-style-type: none"> • HSG is associated with individual ports • HSGs only allow one initiator for host security masking. • To create a host security group, you must specify a port, initiator and a volume. • Every port has a LUN HSG, even if no LUNs are defined for that port. To bind a LUN to a port, edit the host security group and add the desired LUN to a port. • You can also add LUNs to a Mask HSG. To add initiators, you need to create the HSG.

The number of initiators allowed in the HSG depends on the type of storage system:

Table 46 Allowed Initiators in Host Security Groups

Storage System	Allowed Initiators in HSGs
EMC CLARiiON	If All Ports is selected on page one of the wizard for creating HSGs, then the initiator tab is applicable and they can choose 0 or more initiators.
Engenio (LSI)	0 or more initiators
HDS	0 or more initiators
HP XP storage systems using Command View XP Advanced Edition	0 or more initiators
IBM	only one initiator per host security group
Sun 6920	only one initiator per host security group
Sun 6130	0 or more initiators
Symmetrix	one initiator for host security masking

NOTE: For the “Volume Creation and LUN Security” option in Path Provisioning, the All Ports node is not shown because volumes cannot be placed inside HSGs for All Ports.

EMC CLARiiON Storage Systems

Upon the creation of volumes, by default the volume is assigned to one of the two controllers. Even though this volume is mapped to a controller, it is not visible from the outside world by a host. The management server reports this volume as unmapped since it is not visible by a host initiator.

- Volumes can be only on SPA or SPB because CLARiiON is active/passive storage, which means it can have only one active path to a volume. Addition of initiators to any of the processors is listed for both processors.
- The HSG is created on all ports of the processor you select unless you select an initiator that uses a different processor and the initiator does not belong to a host security group. For example, assume you select processor SPA. Then, you select an initiator that belongs to SPB but it does not belong to a host security group. The host security group is created for all ports on SPB.
- HSGs can consist of initiators (WWN) only. You do not need to specify volumes. The initiator is shown in both host security groups SPA and SPB.

- HSGs can consist of volumes (LUNs) only. You do not need to specify initiators.
- When you select an initiator for the host security group, the initiator has to be registered with the CLARiiON storage system.
- You can have more than one initiator in a security group if you have the proper multipathing software installed on that particular host where the initiator is located.

Engenio Storage Systems

- When you create the HSG by using the management server, the HSG appears as a host with its volumes and ports displayed underneath the tree in SANtricity.
- When you create the HSG by using the management server, you cannot specify the controller, but not the port for the HSG.
- An initiator is equivalent to a host port in SANtricity.
- You can have multiple volumes and initiators in the HSG.
- If you create a volume on a host and you have no multipathing on the host, make sure the volume is on the preferred path. You can make sure the volume is on the preferred path by using SANtricity.

EMC Symmetrix Storage Systems

If no LUN security is turned on for an FA port, all volumes assigned to the FC port are visible by hosts that are on the SAN and have been zoned by the SAN. All volumes assigned to the FC port appear in the mapped category.

When you create the HSG on a Symmetrix storage system, you are creating LUN mapping and masking in one step. In the native tools for Symmetrix storage systems, you will not see the HSG you created by using the management server. Instead you will see a volume bound to a port and a masked LUN to a host in the native tools.

- HSG is associated with individual ports
- HSGs only allow one initiator for host security masking.
- To create a host security group, you must specify a port, initiator and a volume.
- Every port has a LUN HSG, even if no LUNs are defined for that port. To bind a LUN to a port, edit the host security group and add the desired LUN to a port.
- You can also add LUNs to a Mask HSG. To add initiators, you need to create the HSG.

HDS Storage Systems

- FC port contains only volumes but no initiators (HBA WWN) assignment, the management server displays these volumes as unmapped since no external host can see these volumes yet.
- You can have 0 to multiple initiators in a host security group
- You can have 0 to multiple volumes in a host security group.
- Every HSG is on one port on the array. You can have host security groups with the same name, as long as they are on different ports.
- Host security groups appear in the native tool for HDS storage systems in the logical view, listed by LDEV and in the physical view, listed by port.

- In the native tool for HDS storage systems, host security groups are referred to as a host security domain.
- When you use the management server to create a host security group, the newly created host security group is displayed by the “name” attribute. The “name” attribute is stored only in the HiCommand database and it is not stored in the HSG itself on the device. This is why the “name” attribute does not appear in the native tools. The name attribute can be 50 characters in length. Hitachi storage arrays contain HSG records on the device. These are identified by the “DisplayName” attribute that is read only. The DisplayName is displayed in the native tools. On some HDS storage systems, an additional “nickname” attribute is available to be set on the HSG and its value is stored on the device itself. The management server uses the “name” attribute instead of the “nickname” attribute because not all HDS storage systems support setting the “nickname” attribute.

HP XP Storage Systems using Command View XP Advanced Edition

- FC port contains only volumes but no initiators (HBA WWN) assignment, the management server displays these volumes as unmapped since no external host can see these volumes yet.
- You can have 0 to multiple initiators in a host security group
- You can have 0 to multiple volumes in a host security group.
- Every HSG is on one port on the array. You can have host security groups with the same name, as long as they are on different ports.
- Host security groups appear in Command View XP Advanced Edition in the logical view, listed by LDEV and in the physical view, listed by port.
- In Command View XP Advanced Edition, host security groups are referred to as a host security domain.

IBM Storage Systems

- To assign a host mode to a host security group, you must modify a property, as described in [“Setting the Host Mode for IBM Storage Systems”](#) on page 281.
- You cannot name a host security group on IBM storage systems. The host security group will be given the name of the initiator you select for the host security group.
- The management server can read the names of HSGs created by the native tool.
- Only one initiator per host security group.
- The volume can be assigned to more than one initiator.
- You can select any number of ports from one to all when creating the HSG.
- You can create the HSG with or without LUNs.
- You can add mapped and unmapped volumes to a host security group, but they should have the same host mode.
- The default is all ports are in the host security group. If no ports are selected, the default is used.

Sun 6920 Storage Systems

- You can have two or more initiators assigned to the same volume, but you can have only one initiator to a host security group.
- You can select any number of ports from one to all when creating the HSG.

- You must create the HSG with one initiator and at least one LUN.
- When the management server creates the HSG, the name of the HSG is the initiator WWN.
- The management server can read the names of HSGs created by the native tool.
- Creation of a new HSG with an initiator that already exists in another HSG adds the volumes to the previously created HSG.
- The management server displays the ports connected to the initiator you selected for creating an HSG. If the initiator is not connected to any ports, no ports are displayed.
- If you select a number of ports, only the selected initiators will be part of the host security group.
- Only the storage ports connected to the initiator in the host security group will be included in the host security group, even if the default is all ports. For example, if an initiator is not connected to a storage system, no ports will be in the host security group.

Sun 6130 Storage Systems

- When you create the HSG by using the management server, the HSG appears as a host with its volumes and ports displayed underneath the tree in SANtricity.
- When you create the HSG by using the management server, you cannot specify the controller, but not the port for the HSG.
- An initiator is equivalent to a host port in the native tool.
- You can have multiple volumes and initiators in the HSG.
- If you create a volume on a host and you have no multipathing on the host, make sure the volume is on the preferred path. You can make sure the volume is on the preferred path by using SANtricity.

Managing Host Security Groups

This section describes the following:

- ["Accessing Information About Host Security Groups"](#) on page 275
- ["Creating Host Security Groups"](#) on page 276
- ["Editing Host Security Groups"](#) on page 279
- ["Deleting Host Security Groups"](#) on page 281

Accessing Information About Host Security Groups

Host Security Groups define which initiators (HBA ports) have access to specified storage volumes. They are associated with a fibre-channel port and contain a list of HBA port initiators and the volumes they can detect.

Keep in mind the following:

- Each type of storage system treats host security groups differently. See ["General Provisioning Issues"](#) on page 282.
- Not all HDS storage systems support host security groups. Refer to the documentation accompanying the HDS storage system.

If you use another product to make provisioning changes, you must perform Discovery Data Collection for the management server to be made aware of these changes.

To access information about host security groups:

1. Click **Tools > Storage Essentials > Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **Storage Systems** tab.
3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to access information about host security groups.
4. Click **Step 3 Host Security Group**.

This page lists the following information about host security groups. You can do the following on this page:

- **View all host security groups** - Click the **All** category in the tree. All the host security groups appear in the right pane.
- **View only host security groups assigned to a certain port** - Click a port in the tree. The host security group assigned to the port appears in the right pane.

This page lists information about host security groups under the following columns:

- **Port** (Displayed if applicable) - Lists the port associated with the host security group.
- **Name** - The name of the host security group.
- **Initiators** - Displays one of the following:
 - The caption for the discovered port if the port has been discovered by the management server (for example Columbia:Adapter0 Port 0). A tool tip over the caption gives the full WWN.
 - The WWN if the port's not been discovered.
- **Volumes** - The volumes in the host security group.
- **Host Mode** - Displays the port settings for your operational environment. The settings for the host mode vary by the model of the HDS storage system. With some hardware, you must select a special host mode on the port for the storage system to enable certain servers and HBAs to "see" the LUNs on the port. Refer to your documentation for the HDS storage system.
- **Host Mode 2** (Displayed if applicable) - Optional settings on the port that describe how the host accesses the port. Multiple options exist. Refer to your documentation for HDS storage system.

You can also create, edit and delete host security groups from this page. See the following topics for more information:

- ["Creating Host Security Groups"](#) on page 276
- ["Editing Host Security Groups"](#) on page 279
- ["Deleting Host Security Groups"](#) on page 281
- ["Setting the Host Mode for IBM Storage Systems"](#) on page 281

Creating Host Security Groups

Host Security Groups define which initiators (HBA ports) have access to specified storage volumes. They are associated with a fibre-channel port and contain a list of HBA port initiators and the volumes they can see.

Keep in mind the following:

- (HDS storage systems) When you use the management server to create a host security group, the newly created host security group is displayed by the “name” attribute. The “name” attribute is stored only in the HiCommand database and it is not stored in the HSG itself on the device. This is why the “name” attribute does not appear in the native tools. The name attribute can be 50 characters in length. Hitachi storage arrays contain HSG records on the device. These are identified by the “DisplayName” attribute that is read only. The DisplayName is displayed in the native tools. On some HDS storage systems, an additional “nickname” attribute is available to be set on the HSG and its value is stored on the device itself. The management server uses the “name” attribute instead of the “nickname” attribute because not all HDS storage systems support setting the “nickname” attribute.
- You cannot use the management server to add a host to a “Host Group”. For example, you cannot have nested host groups.
- In releases previous to build 4.0, host security groups were only supported for HDS storage systems. In this build, host security groups are now available to all storage systems that support provisioning. Each storage system treats host security groups differently, see [“Rules for Creating Host Security Groups”](#) on page 270 for more information.
- Each type of storage system treats host security groups differently. See [“General Provisioning Issues”](#) on page 282.

To create a host security group:

1. Click **Tools** > **Storage Essentials** > **Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **Storage Systems** tab.
3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to access information about host security groups.
4. Click **Step 3 Host Security Group**.
5. Click the **New Host Security Group** button in the upper-right corner of the screen.

Step 1 - Add Details for the Host

1. Type a unique name for the host security group in the **Name** field.

NOTE: You cannot name a host security group on IBM storage systems. The host security group will be given the name of the initiator you select for the host security group.

Keep in mind the following:

- The name must contain 1 to 50 characters. If you enter no characters, you are given the option of using a default name.
- The first and last letter cannot be spaces
- You cannot have the following characters in the name:
 - <
 - >
 - ;

- :
- ,
- |
- /
- *
- ?
- \
- \\
- \t
- \n
- \b

2. (Displayed if applicable) Select the port you want associated with the host security group. This port should contain your LUNs.


NOTE: Each type of storage system handles ports for HSGs differently. See [Table 45](#) on page 270.

3. (HDS storage system) Click the **Options** button to the right of the **Host Mode** field. Select a host mode resembling the port settings for your environment and then click **OK**. If your host mode is not listed, type it in the **Host Mode** field.
4. (HDS storage systems) If your storage system supports a second host mode, type the second host mode in the **Second Host Mode** field.
A second host mode is an optional settings on the port that describe how the host accesses the port (not applicable to all storage systems).
5. (IBM storage systems) You cannot assign the host mode for an IBM storage system in the user interface. You must modify an internal property to set the host mode. See "[Setting the Host Mode for IBM Storage Systems](#)" on page 281 for more information.
6. Click **Next**.

Step 2 - Add Initiators to the Host Security Group

1. To add an initiator to the host security group, click the **Add** button in the upper-right corner.
2. Do one of the following:
 - Type the WWN of the port you want to add to the host security group
 - Select the initiator you want to add to the host security group.


Notice that when the mouse hovers over the port, you are shown additional information, such as the name and WWN of the port on the switch that the host uses.
3. Click the **Add** button at the bottom of the window.
4. When you are done with adding initiators, click the **Close** button.

5. To remove an initiator from the host security group, click the  button. To remove multiple HBA initiators from the list, select the HBA ports you want to remove and then click the **Remove Selected** button.
6. Click **Next**.

Step 3 - Add Volumes to the Host Security Group


1. To add a volume to the host security group, click the **Add** button in the upper-right corner of the window.
2. Select a volume. Then, do one of the following:
 - If you want the unit number to be selected automatically by the server, leave the **Auto-Select** option selected.
 - If you want to choose a unit number, deselect the **Auto-Select** option and enter the unit number in the **Unit Number** field at the top of the window.

NOTE: Keep in mind that for Engenio storage systems LUN numbers cannot be duplicated, and that the management server can use an existing LUN number if the access mode is “No Access” for the created LUN.

3. Documentation available for the assigning LUN number.
4. Click **Add** at the bottom of the window.
The volume is added to host security group.
5. When you are done adding volumes, click the **Close** button.
6. To remove a volume from the host security group, click the  button. To remove multiple volumes from the list, select the volumes you want to remove and then click the **Remove Selected** button.
7. Click **Finish**.

Editing Host Security Groups

To edit a host security group:

1. Click **Tools > Storage Essentials > Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **Storage Systems** tab.
3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to access information about host security groups.
4. Click **Step 3 Host Security Group**.
5. Click the  button corresponding to the host security group you want to edit.


Step 1 - Edit Details for the Host

1. Modify the name of the host security group.
2. Change the port you want associated with the host security group. This port should contain your LUNs.

NOTE: Each type of storage system handles ports for HSGs differently. See [Table 45](#) on page 270.

3. (Only HDS storage system) Click the **Options** button to the right of the **Host Mode** field. Select a host mode resembling the port settings for your environment and then click **OK**. If your host mode is not listed, type it in the **Host Mode** field.
4. (Only HDS storage systems) If your storage system supports a second host mode, type the second host mode in the **Second Host Mode** field.
A second host mode is an optional settings on the port that describe how the host accesses the port (not applicable to all storage systems).
5. (IBM storage systems) You cannot assign the host mode for an IBM storage system in the user interface. You must modify an internal property to set the host mode. See "[Setting the Host Mode for IBM Storage Systems](#)" on page 281 for more information.
6. Click **Next**.

Step 2 - Change the Initiators Assigned to the Host Security Group

1. Change the initiators assigned to the host security group, by doing one or more of the following:
 - **Add an Initiator** - Click the **Add** button. Then, do one of the following:
 - Type the WWN of the port you want to add to the host security group
 - Select the initiator you want to add to the host security group.Notice that when the mouse hovers over the port, you are shown additional information, such as the name and WWN of the port on the switch that the host uses.
Then, click **Add**. Click **Close** to exit the window.
 - **Delete an initiator** - Click the  corresponding to the initiator you want to remove.
 - **Delete multiple initiators** - To remove multiple initiators from the host security group, select the initiators you want to remove and then click **Remove Selected**.

CAUTION: Removing an HBA can cause hosts that are using it to lose access to their storage. This may result in data loss.

2. Click **Next**.

Step 3 - Change the Volumes Assigned to the Host Security Group

IMPORTANT: You cannot delete the default host security group.


1. Change the volumes assigned to the host security group, by doing one or more of the following:
 - **Add a volume** - Click the **Add** button, select a volume and then, do one of the following:
 - If you want the unit number to be selected automatically by the server, leave the **Auto-Select** option selected.

- If you want to choose a unit number, deselect the **Auto-Select** option and enter the unit number in the **Unit Number** field at the top of the window.

Then, click **Add**. Click **Close** to exit the window.

NOTE: Keep in mind that for Engenio storage systems LUN numbers cannot be duplicated, and that the management server can use an existing LUN number if the access mode is “No Access” for the created LUN.

CAUTION: Removing an HBA can cause hosts that are using it to lose access to their storage. This may result in data loss.


- **Remove a volume** - Click  corresponding to the volume you want to remove from the host security group.
- **Remove multiple volumes** - To remove multiple volumes from the host security group, select the volumes you want to remove and then click **Remove Selected**.

2. Click **Finish**.

Deleting Host Security Groups

IMPORTANT: You cannot delete the default host security group.

To delete a host security group:

1. Click **Tools** > **Storage Essentials** > **Provisioning Manager** in HP Systems Insight Manager.
2. In the right pane, click the **Storage Systems** tab.
3. In the right pane, click the **Provision** button corresponding to the storage system in which you want to access information about host security groups.
4. Click **Step 3 Host Security Group**.
5. Click the  button corresponding to the host security group you want to delete.
The host security group is removed.
6. To remove multiple host security groups, select the host security groups you want to remove.
Then click **Delete Selected**.

Setting the Host Mode for IBM Storage Systems

The host mode for an IBM storage system cannot be set in the user interface. You can set the host mode by assigning the host mode to the `smi.ProvisioningIbmEss.hostConnectionProfile` property, as described in the following steps:

1. Click **Options** > **Storage Essentials** > **Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
1. Click **Show Default Properties** at the bottom of the page.

2. Copy `smi.ProvisioningIbmEss.hostConnectionProfile`. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
3. Return to the Advanced page (**Options > Storage Essentials > Manage Product Health**). Then, click **Advanced** in the Disk Space tree).
4. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
5. Assign a host mode to `smi.ProvisioningIbmEss.hostConnectionProfile` in the **Custom Properties** field, as shown in the following example:

```
smi.ProvisioningIbmEss.hostConnectionProfile=AIX
```

where AIX is the host mode.
6. When you are done, click **Save**.
7. Restart the service for the management server for your changes to take effect.
8. While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

General Provisioning Issues

This section describes the following:

- ["Provisioning Can Make a Device Inaccessible"](#) on page 282
- ["Provisioning Does Not Make an Operating System Aware of a Device"](#) on page 282

Provisioning Can Make a Device Inaccessible

Keep in mind that provisioning can break a connection between an array and a host. When you rezone a device, make sure no users or applications are using the device, as it will become unavailable. For example, assume a port for a disk drive is a member of zone set A, which is active. If you make zone set A inactive and this port is not a member of the new active zone set, the disk drive will become unavailable.

Provisioning Does Not Make an Operating System Aware of a Device

When a port in a zone set becomes active, you must take the appropriate steps to make it available to the operating system. For example, assume a port for a disk drive is a member of zone set A. If you make the zone set active, the host will not automatically recognize the disk drive. You will need to configure the operating system, so that it becomes aware of the device. Refer to the documentation that accompanies the operating system for the host.

Provisioning Issues by Vendor

This section describes the following:

- ["Issues Specific to CLARiiON Storage Systems"](#) on page 283
- ["Issues Specific to EMC Symmetrix Storage Systems"](#) on page 283
- ["Issues Specific to HDS Storage Systems"](#) on page 285

- ["Issues Specific to Engenio Storage Systems"](#) on page 287

Issues Specific to CLARiiON Storage Systems

The following issues are specific to CLARiiON® storage systems:

- ["Make the Management Server a Privileged User for CLARiiON"](#) on page 283
- ["EMC Navisphere CLI Is Required"](#) on page 283

Make the Management Server a Privileged User for CLARiiON

Before you can provision a CLARiiON storage system, you must configure it to recognize the management server as a privileged user. This task can be completed by using NaviCli as follows:

```
C:\>navicli -h 192.168.1.249 remoteconfig -setconfig -o -adduser
SYSTEM@hostname
```

- where
 - 192.168.1.249 is the URL of the CLARiiON storage system
 - SYSTEM@hostname is the name of the computer running the management server
- This example tells the CLARiiON storage system at 192.168.1.249 to accept configuration change requests from management server running on the computer named hostname.

EMC Navisphere CLI Is Required

The EMC Navisphere® CLI is required for the management server to communicate with the CLARiiON® storage system. At the time this documentation was created, EMC distributed the Navisphere CLI as part of the EMC Navisphere Software Suite. Contact your EMC representative for more information about obtaining the Navisphere CLI. Distribution rights for the Navisphere CLI belong to EMC. EMC Navisphere CLI Is Required

In Navisphere Manager add one of the following to the privilege user section:

```
SYSTEM@name_of_my_management_server
SYSTEM@IP_of_my_management_server
```

where

- name_of_my_management_server is the DNS name of the computer running the management server software
- IP_of_my_management_server is the IP address of the computer running the management server software

When you use the management server to discover the CLARiiON storage system (**Discovery > Setup > Add Address**), provide the IP address for the CLARiiON storage system and the user name and password used to log into Navisphere.

Issues Specific to EMC Symmetrix Storage Systems

The following issues are specific to EMC Symmetrix Storage Systems:

- ["About Provisioning on EMC Symmetrix Storage Systems"](#) on page 284
- ["LUN Provisioning Is Not Supported with Solutions Enabler 4.3"](#) on page 284

- [“Process Has an Exclusive Lock” Message](#) on page 284
- [“Some EMC Volumes, Their LUNs and LUN Maskings Are Hidden”](#) on page 285

About Provisioning on EMC Symmetrix Storage Systems

EMC ships its Symmetrix storage system with volumes already created. When the software first discovers an EMC Symmetrix storage system, it assumes the devices on the Symmetrix storage system are volumes.

This software refers to the term “device” to define a piece of hardware in the storage network. EMC uses the term “device” to refer to a volume on one of its storage systems. In this section, the term “device” is used in the context of EMC storage systems.

When you use the software's “create a volume” feature, you are assigning the already created volume. If necessary the software will create a meta device, which is a device that is a concatenation of several devices.

The software does not delete the volumes created by EMC. When you use the software's “delete a volume” feature, the software marks the volume as hidden in its repository. These “hidden volumes” are stored in the “Free Device” list.

If you use a device in the “Free Device” list when you create a volume, that device is removed from the “Free Device” list. The device was moved from the list because it is now assigned.

LUN Provisioning Is Not Supported with Solutions Enabler 4.3

LUN mapping and LUN Masking works only when SYMAPI server version 5.0 or later is used. SYMAPI version 4.3 does not allow LUN provisioning, even though it shows existing LUNs. It also does not show existing LUN maskings.

“Process Has an Exclusive Lock” Message

You will receive a message resembling the one shown below if a process has already locked the EMC Symmetrix storage system and you attempt a process that requires a lock on the Symmetrix storage system. The Symmetrix storage system can become locked for many reasons. For example, the storage system becomes locked when it performs LUN mapping, LUN masking or Discovery Data Collection. The Symmetrix storage system may also remain locked after a provisioning operation has failed.

“SYMAPI routine SymDevMaskSessionStart failed with error code 188: The operation failed because another process has an exclusive lock on the local Symmetrix.”

After the management server has detected the lock on the Symmetrix storage system, it tries to access the storage system for 15 minutes and logs the errors.

If you receive the error message, determine if someone is performing an operation that requires a lock, such as LUN mapping, LUN masking or Discovery Data Collection. This also applies even if one of the processes is being used by a third-party product, such as for LUN masking. If so, wait until the process is complete. Only manually remove the lock if you are certain that no other processes are occurring on the storage system. To learn how to remove the lock, refer to the documentation for the Symmetrix storage system.

If a provisioning failure has caused the Symmetrix storage system to remain locked, you are alerted to this situation in Event Monitoring for Storage Essentials and on the Properties tab. You may receive a message resembling the following:

```
Unable to end device masking session. Symmetrix '000001835005700' may be locked.
```

Some EMC Volumes, Their LUNs and LUN Maskings Are Hidden

EMC volumes, their LUNs and LUN maskings for the volumes that play a special role (such as holding device masking information, attached BCVs, gate-keeper devices) are hidden. On the host side, the software shows LUN maskings for these volumes. So you may see a LUN masking, but not its volume residing on the EMC storage system.

Issues Specific to HDS Storage Systems

This section describes issues specific to HDS storage systems.

- ["About Provisioning on HDS Storage Systems"](#) on page 285
- ["Unable to Provision When HDS CruiseControl Is Enabled"](#) on page 286
- ["Increasing the Wait Time for the Management Server"](#) on page 286
- ["Initiator Ports Cannot Be Used for Provisioning"](#) on page 287
- ["The Software Is Unable to Map a Host to Volumes in Different Host Storage Domains of the Same Port on HDS 9900V"](#) on page 287

About Provisioning on HDS Storage Systems

IMPORTANT: The management server does not allow LUSE of different RAID levels.

HDS ships some of its storage systems with volumes already created. When the software first discovers an HDS storage system, it detects the volumes created by HDS. When you use the software's "create a volume" feature, you are assigning the already created volume.

The software does not delete volumes created by HDS. When you use the software's "delete a volume" feature, the software marks the volume as hidden in its repository. These "hidden volumes" are stored in the "Free LDEVs" list.

If you use an LDEV in the "Free LDEVs" list when you create a volume, the LDEV is removed from the "Free LDEVs" list. The LDEVs was moved from the list because it is now assigned.

HDS cannot create a LUSE volume (made up of multiple LDEVs) without mapping it to a target port (that is without creating a LUN). In the software, creating a volume and creating a LUN are two different operations. Therefore, the software keeps the volumes, made from multiple LDEVs in the "Grouped LDEVs" list in the repository. Once these volumes are mapped to the target port and a LUN is created for them, they are removed from the repository and a real LUSE volume is created on the HDS box.

For example, assume you have several 2-GB "Free LDEVs" and you want to create a 4-GB volume. Since the request is larger than an LDEV, two of the "Free LDEVs" will be used for the 4-GB volume.

Unable to Provision When HDS CruiseControl Is Enabled

When HDS CruiseControl is enabled on an HDS array, such as an HDS Lightning 9980V, you are unable to do provisioning. You might also receive the following error message:

An error was encountered during this operation. Some of the operation may have been applied to the storage subsystem. A refresh of the storage subsystem is recommended. "The LDEV is HIHSM reserved; cannot be used in a LUSE".

To use the Provisioning tool, disable HDS CruiseControl. Refer to the HDS CruiseControl product documentation for more information.

Increasing the Wait Time for the Management Server

By default, the management server waits 30 minutes for a response from HiCommand Device Manager after sending a command. If the management server does not receive a response from HiCommand Device Manager after 30 minutes, the management server assumes the command did not go through. It then tries to contact HiCommand Device Manager again while the previous command is still proceeding.

For example, assume you initiated a command to add the HSG. The management server waits 30 minutes for a response. If it receives no response during the time, it sends another "Add" command and waits 30 minutes for a response. If no response is received, it sends another command. It continues this pattern until it receives a response from HiCommand Device Manager. As a result, multiple HSGs with the same configuration of LUNs and Worldwide Names (WWNs) are created until an unexpected error occurs or you stop the AppStorManager service.

The management server has a similar behavior when a delete command is initiated. It sends the delete command to HiCommand Device Manager, waits 30 minutes, and sends another delete command if it receives no response. The commands try to delete the same HSG, but the target HSG is deleted after the first command is completed. The second command returns an error.

If you need more time for HiCommand Device Manager to respond, you can increase the amount of time the management server waits, by modifying the `cimom.provider.hds.timeout` property as described in the following steps:

1. Select **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree in the management server.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
`cimom.provider.hds.timeout=1800000`
4. Return to the Advanced page (**Options > Storage Essentials > Manage Product Health**). Then, click **Advanced** in the Disk Space tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make your changes in the **Custom Properties** field. Change the value of the `cimom.provider.hds.timeout` property. The value is in milliseconds. For example, assume

you want the management server to wait an hour. You would assign 3600000 to `cimom.provider.hds.timeout`, since 3600000 milliseconds is one hour.

```
cimom.provider.hds.timeout=3600000
```

7. When you are done, click **Save**.

8. Restart the service for the management server for your changes to take effect.

While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Initiator Ports Cannot Be Used for Provisioning

Ports designated as an Initiator on a storage system belonging to the HDS Freedom Storage™ Lightning 9900™ Series or Freedom Storage Lightning 9900V Series cannot be used for provisioning. If you select one of these ports, you receive a message saying that provisioning failed because the HiCommand Database was not refreshed. The management server does not support provisioning for ports designated as Initiators on these storage systems.

The Software Is Unable to Map a Host to Volumes in Different Host Storage Domains of the Same Port on HDS 9900V

On HDS 9900V storage systems, if a host is already mapped to a volume and you try to map the same host to a volume in another host storage domain, corresponding to the same port, it will fail. However for HDS 9900, the host can be mapped to a volume in another host storage domain corresponding to the same port.

A Default LUN Number Is Used Instead of a User-Specified One

When you create a LUN, the user-specified LUN is ignored. The LUN is created with the next available default number.

Issues Specific to Engenio Storage Systems

This section describes the following for Engenio:

- "Creating and Deleting Storage Pools" on page 287
- "Creating and Deleting Storage Volumes" on page 287
- "Creating a LUN" on page 288

Creating and Deleting Storage Pools

For Engenio, a storage pool is the same as a volume group. Create at least one storage pool before provisioning a volume. When you delete a storage pool on an Engenio storage system, all the volumes for the volume group are deleted, including the placeholder volume.

Creating and Deleting Storage Volumes

Keep in mind the following when you create a volume on an Engenio storage system:

- Make sure you select a volume group that can accommodate the requested size for the new volume.
- You can create volumes from existing free extent areas within a volume group.

- The volume capabilities, their cache read ahead multiplier and segment size are shown in the following table:

Table 47 Volume Usage

Volume Capability	Cache Read Ahead Multiplier	Segment Size
File and Default	1	64 KB
Database	0	64 KB
Multimedia	8	128 KB
Custom	0 to 65,535	8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB

- No volume-to-LUN masking is done by default.
- The volume is assigned to a Host not a Host Group. Host and Host Groups are not viewable in the software, and the LUN number may change when it is mapped to a host if that host already has a volume with the LUN number.

IMPORTANT: The management server creates a placeholder volume when a storage pool is created. This placeholder volume is not viewable in the management server, but it might be viewable in other storage tools. Do not delete this placeholder volume.

Creating a LUN

After a volume is assigned to a storage system port, the management server lists the storage system port assigned to each port on the controller. The LUN information is shown on four lines (For two ports on two controllers), one for each port on each controller. The creation of a LUN maps the volume to all the available ports on an Engenio storage system.

10 Path Provisioning

IMPORTANT: Depending on your license, Path Provisioning may not be available. See the “List of Features” to determine if you have access to Path Provisioning. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).

This chapter describes the following:

- “About Path Provisioning” on page 289
- “How Path Provisioning Works” on page 290
- “How to Use Path Provisioning” on page 291
- “System Actions” on page 292
- “Adding a Host” on page 311
- “Creating a Host Security Group” on page 311
- “Scheduling Provisioning Jobs” on page 313
- “Executing Provisioning Jobs” on page 314
- “Monitoring Provisioning Jobs” on page 315
- “Deleting Multiple Jobs” on page 315
- “Naming Conventions for Zones” on page 316
- “Using Multipathing with Path Provisioning” on page 316
- “Customizing Path Provisioning” on page 317

About Path Provisioning

Path Provisioning lets you schedule provisioning tasks to take place at a later time when the network traffic is light. For example, you could use Path Provisioning to schedule multiple provisioning tasks to take place at 1 a.m. When you come in the later that morning, you can view the status of the provisioning tasks.



You can also use Path Provisioning to identify host/storage dependencies so you can make informed decisions when deciding where new volumes, zones or LUN security is needed.

To view the latest provisioning information in Path Provisioning, click the **Refresh** button. The **Refresh** button updates the following:

- The Path Provisioning screen with changes made in Provisioning Manager. For example, assume you used the wizards in Provisioning Manager to create a host security group. When you access Path Provisioning, your changes are not shown.
- The Path Provisioning screen with changes from executed jobs. After a job is executed in Path Provisioning, the Path Provisioning screen is not updated until you click the **Refresh** button or you exit and re-enter Path Provisioning.
- Other parts of the product, such as:
 - Application Viewer

- Capacity Manager
- Performance Manager
- Provisioning Manager
- System Manager

Keep in mind the following:

- Path Provisioning runs within a Java applet. If you receive “out of memory” messages when you view Path Provisioning, you may need to increase the amount of memory assigned to the Java plug-in on the client computer.
- If you select a direct attached host, the storage ports appear in the LUN pane with a  next to them, indicating they are unreachable. You can still select these storage ports that have a  next to their icon and schedule the job. These storage ports are shown as unreachable in the user interface because the user interface uses switches to display the association between a host and a storage system. If the management server cannot detect a switch, as with a DAS connection, the user interface assumes the storage ports are unreachable.
- If the management server there is no switch, the user interface assumes they are not connected.
- Path Provisioning displays all elements the management server detects, even those that do not belong to a user's organization.
- Ports designated as an Initiator on a storage system belonging to the HDS Freedom Storage™ Lightning 9900™ Series or Freedom Storage Lightning 9900V Series cannot be used for provisioning. If you select one of these ports, you receive a message saying that provisioning failed because the HiCommand Database was not refreshed. The management server does not support provisioning for ports designated as Initiators on these storage systems.
- The HBAs displayed may not have a connection to the selected storage system. This provides flexibility. You can select a disconnected HBA for a job to take place when it will be connected to the storage system.

Keep in mind the following for McDATA and Connectrix switches:

- When McDATA or Connectrix switches are discovered through a proxy by using SNMP, you cannot view or perform any provisioning operations for those switches. For example, you cannot view zone sets, zones, and/or zone aliases.
- When McDATA or Connectrix switches are discovered by their IP address by using SNMP, you can only view the active zone set and its members. You cannot create, modify, and/or delete zone sets or its members.
- Zone aliases are not supported for McDATA or Connectrix switches.
- You can view zones, zone sets and zone aliases on a Cisco switch; however, you cannot use the management server to create, modify or remove them from a Cisco switch.
- Path Provisioning looks for the names of the active zone set and of the active zones and all of their saved counterparts in the zoning library in EFC Manager. The provisioning job only occurs if those names match.

How Path Provisioning Works

When you select a storage system in Path Provisioning, the management server displays all the information relevant for provisioning with respect to the storage system selected. This information includes:


- Mapped, unmapped and unmasked (mapped to one or more storage system ports but not associated to any host initiator port) volumes of the storage system,
- Already masked LUNs of the storage system and the front end ports of the storage system (possible candidates for LUN mapping step)
- Hosts that are reachable (hosts belonging to the same fabric as the storage system) along with their HBA and host initiator ports. Single multipathing functionality is supported.
- Existing zones to which the ports of the storage system belong are displayed. If the Host and the Storage System belong to multiple fabrics, zones of all those fabrics are displayed.

The above information is displayed in several different panes (storage system, host, volume, LUN and zone panes). You can select relevant information from each of these panes (such as, which volumes to be used for the LUN Mapping Task). Each of these selections potentially can create a provisioning task to be performed. You can create a job that contains one of the following:

- A single task of one type (such as, Map LUN)
- Set of tasks of the same type (such as, multiple LUN Mapping)
- Multiple LUN Masking tasks or a job consisting of several tasks of different types (such as, Map LUN, zone required ports, mask LUN, which go together as a combination).
- You can also create multiple jobs each consisting of multiple tasks of different types.

The jobs can be executed immediately or can be scheduled to start at a later time. Jobs with all its required details (such as, parameters to invoke Provider Service methods) are stored in a job queue. The scheduler, at the scheduled time, retrieves the job from the job queue and performs the tasks of the job using the details stored in the job.

The status of each of the job can be one of the following, and it is displayed in the State column in the Provision Job section of the screen (lower pane):

- **Create** - The job has been created, but it will not be executed. The job cannot be viewed by others and it is deleted when the Web browser is closed. See "[Scheduling Provisioning Jobs](#)" on page 313 for information about changing the state of the job from "create" to "scheduled".
- **Scheduled** - The job has been tasked to execute at a specified time and date. Jobs are assigned a scheduled state after you select the job and click the **Execute Job** button ( button).
- **Started** - The job has started. You cannot delete a job once it has started.
- **Finished** - The job has finished.

How to Use Path Provisioning

All the provisioning tasks are centralized on one screen. You select the provisioning actions you want to take by selecting a system action from the System Action combo-box in the upper-left corner. The following options are available:


- **Volume Creations + LUN Security + Zone Operation**- First you create a meta volume. Second, you map a meta volume to a Fibre Channel port and host HBAs (HSG). Third, you create a zone. See "[Volume Creation, LUN Security, and Zone Operation](#)" on page 293 for more information.
- **Meta Volume Creation** - Creates a meta volume. See "[Meta Volume Creation](#)" on page 298 for more information.
- **LUN Security** - Map meta volume to Fibre Channel port and host HBAs (HSG). See the topic "[LUN Security](#)" on page 299 for more information.
- **Zone Operation** - Lets you perform just a zone operation. See "[Zone Operation](#)" on page 304 for more information.
- **Volume Creation + LUN Security** - First you create a meta volume. Second you map a meta volume to a Fibre Channel port). See "[Volume Creation and LUN Security](#)" on page 308 for more information.

After you select a system action, complete the steps in the various panes. If a step is disabled after all data has been loaded, that step is not required.

When you have completed your steps, schedule the task as described in the topic, "[Scheduling Provisioning Jobs](#)" on page 313.

The following table provides an overview of the steps required for Path Provisioning.

Table 48 Overview for Path Provisioning

Step	Description	Where to Find Additional Information
1	Select a system action.	See the topics:
2	Complete the steps.	<ul style="list-style-type: none"> • "Volume Creation, LUN Security, and Zone Operation" on page 293 • "Meta Volume Creation" on page 298 • "LUN Security" on page 299 • "Zone Operation" on page 304 • "Volume Creation and LUN Security" on page 308
3	(optional) Schedule a provisioning job if you want it to take place at a later time.	See " Scheduling Provisioning Jobs " on page 313.
4	Execute the job. The job does not run until the Execute Job () button is clicked. Once this button is clicked the job is saved in the management server database.	See " Executing Provisioning Jobs " on page 314.

System Actions



This section describes system actions for the following:

- “Volume Creation, LUN Security, and Zone Operation” on page 293
- “Meta Volume Creation” on page 298
- “LUN Security” on page 299
- “Zone Operation” on page 304
- “Volume Creation and LUN Security” on page 308
- “Providing a LUN Number” on page 310
- “Adding a Host” on page 311

Volume Creation, LUN Security, and Zone Operation

You can use Path Provisioning to create meta volumes. Then, map that meta volume to a Fibre Channel port and host HBA (HSG). Finally, you could designate that volume to appear in a pre-existing zone or create your own.

Keep in mind the following:

- If you have options still selected from a previous job, just clear the options you do not want in your next job. For example, assume you created a volume. Now you want to create a new volume on the same host used previously. Just clear the Volume pane. To clear a pane, click the  button.
- If you want to clear all the steps, except for the Step 1 (storage systems), select another option from the System Action combo-box.
- (HDS only) Before you can create a volume. You must delete some unmapped LDEVs using the standard provisioning tool. To delete LDEVs, click **Provisioning** (). Then, click the Storage Systems tab, then the **Provision** button for the storage system. Then, click **Step 2 Volume**. Click the **Unmapped Volumes** folder in the left pane. Select the desired number of LDEVs for the LUSE volume then click the **Delete Selected Volumes** button. Take note of the array group from which you deleted the LDEVs. You need this information to create the LUSE volume.

To access Path Provisioning:

1. Click **Tools > Storage Essentials > Provisioning Manager**.
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Select the following from the **System Action** combo-box: Volume Creation + LUN Security + Zone Operation

Step 1 - Select Storage System

1. Wait for the management server to load the storage systems into the Storage System panel.
2. Select a storage system on which you want to create the metavolume.

NOTE: The S heading in the Storage Systems pane means that only a single selection is allowed.

(HDS only) Select the storage system from which you want to create the LUSE volume.

3. Click the **Step 1** button underneath the pane.

Information about the selected storage system, such as its model and IP address is displayed below the Storage System pane. The Host pane becomes populated. Notice in the following pane that some hosts have a red X over their icon. This means the host is not accessible.

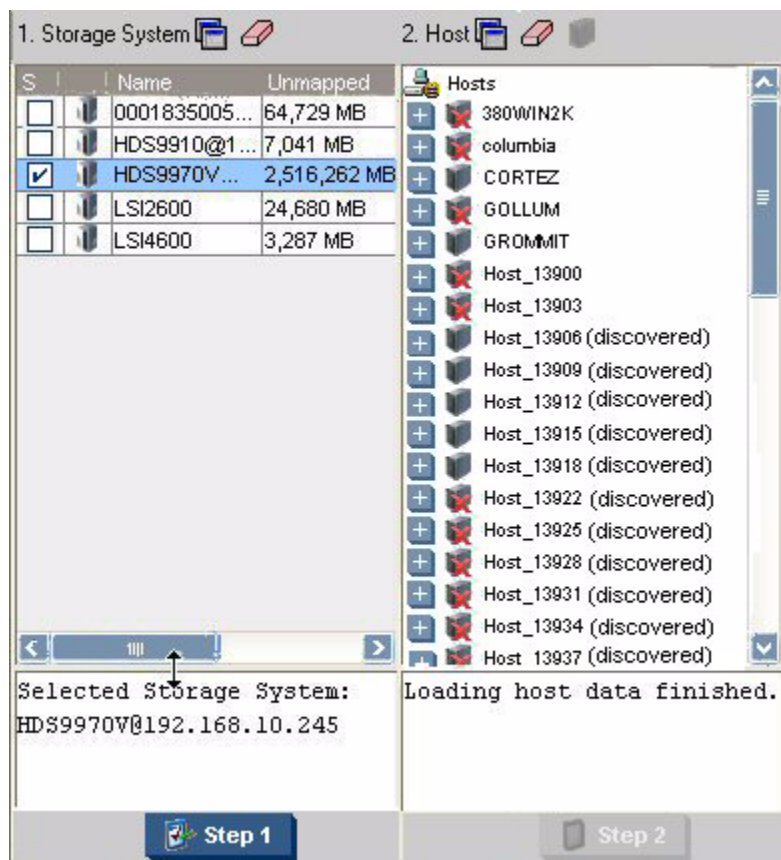


Figure 53 Selecting a Storage System




Step 2 - Select a Host


1. Wait for all data to be loaded. You can tell if all data has been loaded when you see the following messages:
 - Loading host data finished.
 - Loading volume data finished.
 - Loading HSG data finished.
 - Loading zone data finished.

The Step 2 button appears disabled until data has been loaded

2. Select a host that is reachable.

Keep in mind the following:

- To create a provisioning job for multipathing, click the  button above the Host pane. Then, select the option for multipathing. When you select this option, you must select the same host for both paths. See "[Host Customize Dialog](#)" on page 317 for more information.
- To automatically create a zone if a zone does not meet a preset criteria, click the  button above the Zone pane. Then, select the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. Select one of the options listed. See "[Zone Customize Dialog](#)" on page 318 for more information about setting the criteria. After you set the criteria, select a host and click **Step 2**. If the management server finds a zone that meets the criteria, it selects the zone in the Zone pane. If the management server does not find a zone that meets the criteria, it asks for a zone name. See "[Naming Conventions for Zones](#)" on page 316 for more information about the naming requirements for a zone. After you type a zone name, the new zone is displayed in the Zone pane, but it has not been created yet. It is created when the job runs.
- To configure zoning manually, click the  button above the Zone pane. Then, deselect the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. In the zone pane, you must select a zone or create one manually.

If you see the  icon next to a host, it means the host is unreachable.

3. Click the **Step 2** button.

Information about the selected port, such as its Worldwide Name, is displayed below the Host pane. The volumes for that host are displayed in the Volume pane.

4. Select a volume in the Volume pane and then click the **Step 3** button.

Step 3 - Select a Volume

To select a volume:



1. In the Volume pane select mapped and unmapped volumes. You can select multiple volumes on Windows computers by pressing the Ctrl key as you select the volumes.
 - **Mapped** - There are two types of mapped volumes:
 - **Masked** - The volume is exposed to the storage port and to the host.
 - **Unmasked** - The volume is exposed to the storage port, but not to the host.
 - **Unmapped** - The volume is not exposed to the storage port. The management server puts all unmapped volumes in this category when it first discovers an array.
 - **Free Extents** - Available free extents that can be used to create a meta volume. You can create meta volumes on EMC Symmetrix and LUSE on HDS storage systems. To create a meta volume or LUSE, select multiple free extents under the Free Extents node in the Volume pane. Select multiple LDEVs from the Free Extents menu by holding down the shift key on your keyboard and selecting free LDEVs. When you select free extents, they must of the

same type. For example, on Symmetrix, you cannot select a mirrored volume and a BCV (business continuous volume) to create a meta volume.

When you first discover a storage system, no free extents are displayed. This is because the management server puts all unmapped volumes into the “unmapped” category for the list of volumes by default. To move a volume to the free extent node, delete the unmapped volume. When the volume is deleted, it is moved to the free extents node. The free extents category is used internally by the management server. See “[Deleting a Storage Volume](#)” on page 268 for more information.

IMPORTANT: Make sure the free extents you select are not being used. Data on the free extents becomes unusable when a meta volume is created.

Keep in mind the following:

- You can narrow the type of volumes displayed in the Volumes pane by using the Volume Customize Dialog () located above the Volumes pane. See “[Host Customize Dialog](#)” on page 317 for more information.
- If the LUN has already been selected and Step 4 is clicked, skip this step or click the  button.

2. Click **Step 3**.


3. If you are asked to specify a LUN number, provide a LUN for each volume displayed. See “[Providing a LUN Number](#)” on page 310 for information about numbering LUNs.

Step 4 - Select a Host Security Group


1. Select a host security group. See “[Creating a Host Security Group](#)” on page 311 for information on how to create a host security group. See “[General Provisioning Issues](#)” on page 282 for information on how your storage system handles host security groups.


2. Click **Step 4**.

Step 5 - Select a Zone

NOTE: If the zone has already been selected and Step 5 is clicked, skip this step or click the  button to clear the selection.

If you do not have a zone selected or you do not have one created yet, most likely the option,

Automatically Configure Zoning, is not selected in the Zone Customize Dialog (). The management server assumes you want to select a pre-existing zone or create one manually when you deselect the **Automatically Configure Zoning** option.

- **To reuse a zone** - Select a zone in the Zone pane and then click **Step 5**. Expand a fabric node to view its zones.
- **To create a zone** - Select a fabric in the zone pane, click the  button, type a zone name. See the following topic for more information.





If you see the  icon, the fabric cannot be reached.


Table 49 Zone Icons

Icon	Description
	Its location on the page determines its use: <ul style="list-style-type: none">• Displayed above Zone pane - Button for creating zones.• In Zone pane - Icon for a zone.
	Zone Alias
	Port

IMPORTANT: (McDATA switches only) Path Provisioning looks for the names of the active zone set and of the active zones and all of their saved counterparts in the zoning library in EFC Manager. The provisioning job only occurs if those names match.

Creating a Zone

To create a zone:

1. Select a fabric in the zone pane.
2. Click the  button located above the Zone pane.
3. Type a zone name in the dialog box. The following restrictions apply:

Naming Conventions for Brocade Switches:

- The name must contain 1 to 64 characters.
- The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).
- The name is case sensitive. For example, "Zone1" and "zone1" are different zones.
- You cannot create a zone with the same name as an existing zone, zone alias or zone set. For example, if you create a zone named "new", you cannot give a zone, zone alias, or zone set the same name.
- The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign (\$).

Naming Conventions for McDATA and Connectrix Switches:

- The name can have a maximum of 64 characters.
- The first character of a zone name must be a letter (A-Z, AZ).
- A zone name cannot contain spaces.
- Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).
- All names must be unique and may not differ by case. For example, myzone and MyZone are considered to be the same zone.


4. Click **OK**.

The new zone is added to the Zone pane.

5. Click the **Create Job** button.



The job is added to the bottom pane.

6. Do one of the following:

- If you want the job to execute now, click the Execute Job () button
- If you want the job to execute at a later time, schedule the job as described in the topic, "[Scheduling Provisioning Jobs](#)" on page 313.

Meta Volume Creation

Keep in mind the following when creating meta volumes:

- If you have options still selected from a previous job, just clear the options you do not want in your next job. For example, assume you created a volume. Now you want to create a new volume on the same host used previously. Just clear the Volume pane. To clear a pane, click the  button.
- If you want to clear all the steps, except for the Step 1 (storage systems), select another option from the System Action combo-box.
- (HDS only) Before you can create a volume. You must delete some unmapped LDEVs using the standard provisioning tool. To delete LDEVs, click **Provisioning** (). Then, click the Storage Systems tab, then the **Provision** button for the storage system. Then, click **Step 2 Volume**. Click the **Unmapped Volumes** folder in the left pane. Select the desired number of LDEVs for the LUSE volume then click the **Delete Selected Volumes** button. Take note of the array group from which you deleted the LDEVs. You need this information to create the LUSE volume.

You can use Path Provisioning to create meta volumes, as described in the following steps:

1. Click **Tools > Storage Essentials > Provisioning Manager**.
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Select the following from the System Action combo-box: **Meta Volume Creation**
4. Wait for the management server to load the storage systems into the Storage System panel.
5. Select the storage system on which you want to create the metavolume.
(HDS only) Select the storage system from which you want to create the LUSE volume.



NOTE: The S heading in the Storage Systems pane means that only a single selection is allowed.

6. Click the **Step 1** button underneath the pane.
Information about the selected storage system, such as its model and IP address is displayed below the Storage System pane.
7. In the Volume pane, select multiple free extents under the Free Extents node in the Volume pane.
The meta volume containing the selected free extents is created when the job runs.

- **Mapped** - There are two types of mapped volumes:
 - **Masked** - The volume is exposed to the storage port and to the host.
 - **Unmasked** - The volume is exposed to the storage port, but not to the host.
- **Unmapped** – The volume is not exposed to the storage port. The management server puts all unmapped volumes in this category when it first discovers an array.
- **Free Extents** – Available free extents that can be used to create a meta volume. You can create meta volumes on EMC Symmetrix and LUSE on HDS storage systems. To create a meta volume or LUSE, select multiple free extents under the Free Extents node in the Volume pane. Select multiple LDEVs from the Free Extents menu by holding down the shift key on your keyboard and selecting free LDEVs. When you select free extents, they must be of the same type. For example, on Symmetrix, you cannot select a mirrored volume and a BCV (business continuous volume) to create a meta volume.

When you first discover a storage system, no free extents are displayed. This is because the management server puts all unmapped volumes into the “unmapped” category for the list of volumes by default. To move a volume to the free extent node, delete the unmapped volume. When the volume is deleted, it is moved to the free extents node. The free extents category is used internally by the management server. See [“Deleting a Storage Volume”](#) on page 268 for more information.

Keep in mind the following:


- Make sure the free extents you select are not being used. Data on the free extents becomes unusable when a meta volume is created.
- You can narrow the type of volumes displayed in the Volumes pane by using the Volume Customize Dialog () located above the Volumes pane. See [“Host Customize Dialog”](#) on page 317.
- If the LUN has already been selected and Step 4 is clicked, skip this step or click the  button.

8. Click **Step 3**.

9. Click the **Create Job** button.


The job is added to the bottom pane.

10. Do one of the following:

- If you want the job to execute now, select the job and then click the Execute Job () button.
- If you want the job to execute at a later time, schedule the job as described in the topic, [“Scheduling Provisioning Jobs”](#) on page 313.

LUN Security

Keep in mind the following:

- If you have options still selected from a previous job, just clear the options you do not want in your next job. For example, assume you just mapped a volume to a port. Now you want to map a different volume on the same host to another port. Just clear the Volume and LUN panes. To clear a pane, click the  button.

- If you want to clear all the steps, except for the Step 1 (storage systems), select another option from the System Action combo-box.

You can use Path Provisioning to designate subsystem LUN security, as described in the following steps:

1. Click **Tools > Storage Essentials > Provisioning Manager**.
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Select the following from the System Action combo-box: LUN Security

Step 1 - Select Storage System

1. Wait for the management server to load the storage systems into the Storage System panel.
2. Select the storage system on which you want to create the metavolume.

NOTE: The S heading in the Storage Systems pane means that only a single selection is allowed.

3. Click the **Step 1** button underneath the pane.

Information about the selected storage system, such as its model and IP address is displayed below the Storage System pane. The Host pane becomes populated. Notice in the following pane that some hosts have a red X over their icon. This means the host is not accessible.

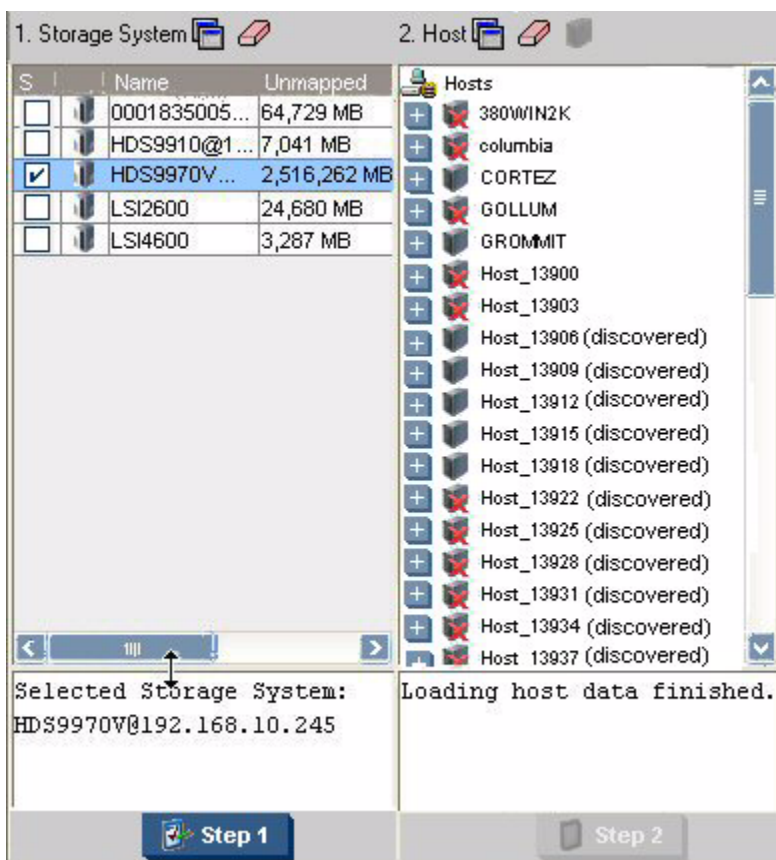



Figure 54 Selecting a Storage System




Step 2 - Select a Host


- Wait for all data to be loaded. You can tell if all data has been loaded when you see the following messages:
 - Loading host data finished.
 - Loading volume data finished.
 - Loading HSG data finished.

The Step 2 button appears disabled until data has been loaded

- Do one of the following:
 - Select a host that is reachable.
 - Add a host that is not currently connected to the network by clicking the  button. See ["Adding a Host"](#) on page 302.

Keep in mind the following:

- To create a provisioning job for multipathing, click the  button above the Host pane. Then, select the option for multipathing. When you select this option, you must select the same host for both paths. See “[Host Customize Dialog](#)” on page 317 for more information.
- To automatically create a zone if a zone does not meet a preset criteria, click the  button above the Zone pane. Then, select the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. Select one of the options listed. See “[Zone Customize Dialog](#)” on page 318 for more information about setting the criteria. After you set the criteria, select a host and click **Step 2**. If the management server finds a zone that meets the criteria, it selects the zone in the Zone pane. If the management server does not find a zone that meets the criteria, it asks for a zone name. See “[Naming Conventions for Zones](#)” on page 316 for more information about the naming requirements for a zone. After you type a zone name, the new zone is displayed in the Zone pane, but it has not been created yet. It is created when the job runs.
- To configure zoning manually, click the  button above the Zone pane. Then, deselect the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. In the zone pane, you must select a zone or create one manually.

If you see the  icon next to a host, it means the host is unreachable.



3. Click the **Step 2** button.

Information about the selected port, such as its Worldwide Name, is displayed below the Host pane. The volumes for that host are displayed in the Volume pane.

Adding a Host

The management server lets you add hosts that are not currently connected to the network. While you are creating a job, add the host.

IMPORTANT: Make sure the added host is physically connected to the network before the scheduled job runs.

1. Click the  button.
2. Type the name for the host in the **Host Name** field.
3. Type the port name of the host in the **Port WWN** field.
4. Click the **Add** button.
5. Repeat Steps 2 and 3 for multiple ports.
6. If you want to remove the host, click the  button.
7. When you are done with your changes, click **OK**.
The host is added to the list of hosts.
8. Physically connect the host to the network before the job runs.

Step 3 - Select a Volume

To select a volume:



1. In the Volume pane select mapped and unmapped volumes. You can select multiple volumes on Windows computers by pressing the Ctrl key as you select the volumes.
 - **Mapped** - There are two types of mapped volumes:
 - **Masked** - The volume is exposed to the storage port and to the host.
 - **Unmasked** - The volume is exposed to the storage port, but not to the host.
 - **Unmapped** – The volume is not exposed to the storage port. The management server puts all unmapped volumes in this category when it first discovers an array.
 - **Free Extents** – Available free extents that can be used to create a meta volume. You can create meta volumes on EMC Symmetrix and LUSE on HDS storage systems. To create a meta volume or LUSE, select multiple free extents under the Free Extents node in the Volume pane. When you select free extents, they must of the same type. For example, on Symmetrix, you cannot select a mirrored volume and a BCV (business continuous volume) to create a meta volume.

When you first discover a storage system, no free extents are displayed. This is because the management server puts all unmapped volumes into the “unmapped” category for the list of volumes by default. To move a volume to the free extent node, delete the unmapped volume. When the volume is deleted, it is moved to the free extents node. The free extents category is used internally by the management server. See [“Deleting a Storage Volume”](#) on page 268 for more information.

IMPORTANT: Make sure the free extents you select are not being used. Data on the free extents becomes unusable when a meta volume is created.

2. Click **Step 3**.
3. If you are asked to specify a LUN number, provide a LUN for each volume displayed. See [“Providing a LUN Number”](#) on page 310 for information about numbering LUNs.

Keep in mind the following:


 - You can narrow the type of volumes displayed in the Volumes pane by using the Volume Customize Dialog () located above the Volumes pane. See [“Host Customize Dialog”](#) on page 317.
 - If the LUN has already been selected and Step 4 is clicked, skip this step or click the  button.

Step 4 - Select a Host Security Group

1. Select a host security group. See [“Creating a Host Security Group”](#) on page 311 for information on how to create a host security group. See [“General Provisioning Issues”](#) on page 282 for information on how your storage system handles host security groups.
2. Click **Step 4**.
3. Click the **Create Job** button.


The job is added to the bottom pane.

4. Do one of the following:


- If you want the job to execute now, click the Execute Job () button
- If you want the job to execute at a later time, schedule the job as described in the topic, ["Scheduling Provisioning Jobs"](#) on page 313.

Zone Operation

Keep in mind the following:

- If you have options still selected from a previous job, just clear the options you do not want in your next job. For example, assume you created a zone. Now you want to create a new zone that includes the same host used previously. Just clear the Zone pane. To clear a pane, click the  button.
- If you want to clear all the steps, except for the Step 1 (storage systems), select another option from the System Action combo-box.

You can use Path Provisioning to perform just zoning operations, as described in the following steps:

1. Click **Provisioning** () .
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Select the following from the System Action combo-box: Zone Operation

Step 1 - Select Storage System

1. Wait for the management server to load the storage systems into the Storage System panel.
2. Select a storage system.

NOTE: The S heading in the Storage Systems pane means that only a single selection is allowed.

3. Click the **Step 1** button underneath the pane.

Information about the selected storage system, such as its model and IP address is displayed below the Storage System pane. The Host pane becomes populated. Notice in the following pane that some hosts have a red X over their icon. This means the host is not accessible.

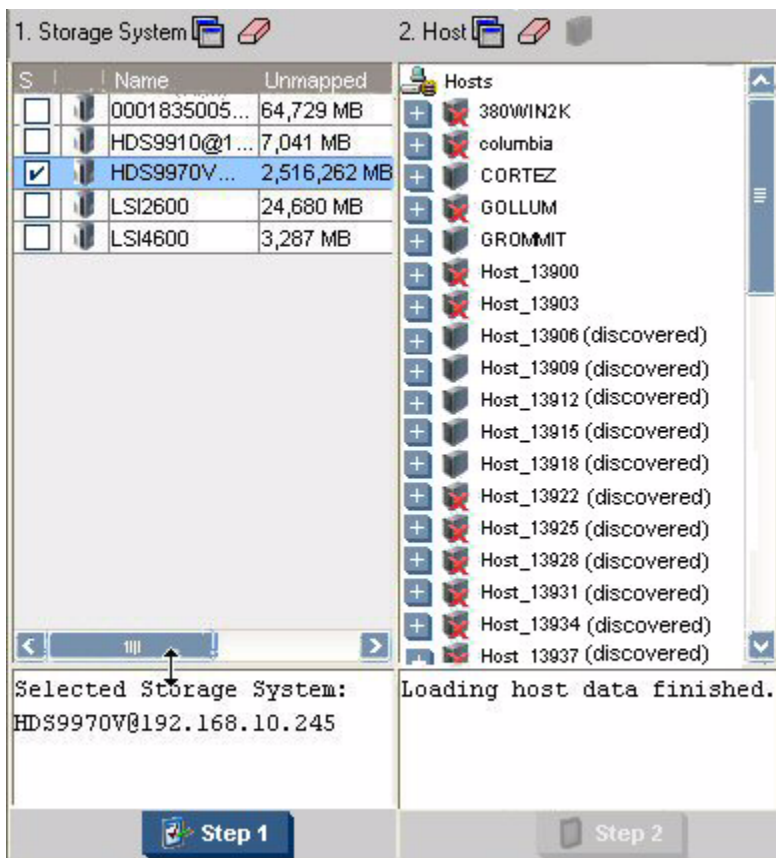


Figure 55 Selecting a Storage System

Step 2 - Select a Host


1. Wait for all data to be loaded. You can tell if all data has been loaded when you see the following messages:



- Loading host data finished.
- Loading HSG data finished.
- Loading zone data finished.


The Step 2 button appears disabled until data has been loaded

2. Select a host that is reachable.

Keep in mind the following:

- To create a provisioning job for multipathing, click the  button above the Host pane. Then, select the option for multipathing. When you select this option, you must select the same host for both paths. See "[Host Customize Dialog](#)" on page 317 for more information.

- To automatically create a zone if a zone does not meet a preset criteria, click the  button above the Zone pane. Then, select the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. Select one of the options listed. See ["Zone Customize Dialog"](#) on page 318 for more information about setting the criteria. After you set the criteria, select a host and click **Step 2**. If the management server finds a zone that meets the criteria, it selects the zone in the Zone pane. If the management server does not find a zone that meets the criteria, it asks for a zone name. See ["Naming Conventions for Zones"](#) on page 316 for more information about the naming requirements for a zone. After you type a zone name, the new zone is displayed in the Zone pane, but it has not been created yet. It is created when the job runs.
- To configure zoning manually, click the  button above the Zone pane. Then, deselect the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. In the zone pane, you must select a zone or create one manually.


If you see the  icon next to a host, it means the host is unreachable.


3. Click the **Step 2** button.


Information about the selected port, such as its Worldwide Name, is displayed below the Host pane. The volumes for that host are displayed in the Volume pane.

- 4.** Select a host security group in the LUN pane. See ["Creating a Host Security Group"](#) on page 311 for information on how to create a host security group. See ["General Provisioning Issues"](#) on page 282 for information on how your storage system handles host security groups.
- 5.** Click **Step 4**.

Step 3 - Select a Zone

NOTE: If the zone has already been selected and Step 5 is clicked, skip this step or click the  button to clear the selection.

If you do not have a zone selected or you do not have one created yet, most likely the option, **Automatically Configure Zoning**, is not selected in the Zone Customize Dialog (). The management server assumes you want to select a pre-existing zone or create one manually when you deselect the **Automatically Configure Zoning** option.

- **To reuse a zone** - Select a zone in the Zone pane and then click **Step 5**. Expand a fabric node to view its zones.
- **To create a zone** - Select a fabric in the zone pane, click the  button, type a zone name. See the following topic for more information.





If you see the  icon, the fabric cannot be reached.


Table 50 Zone Icons

Icon	Description
	Its location on the page determines its use: <ul style="list-style-type: none">• Displayed above Zone pane - Button for creating zones.• In Zone pane - Icon for a zone.
	Zone Alias
	Port

IMPORTANT: (McDATA switches only) Path Provisioning looks for the names of the active zone set and of the active zones and all of their saved counterparts in the zoning library in EFC Manager. The provisioning job only occurs if those names match.

Creating a Zone

To create a zone:

1. Select a fabric in the zone pane.
2. Click the  button located above the Zone pane.
3. Type a zone name in the dialog box. The following restrictions apply:

Naming Conventions for Brocade Switches:

- The name must contain 1 to 64 characters.
- The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).
- The name is case sensitive. For example, "Zone1" and "zone1" are different zones.
- You cannot create a zone with the same name as an existing zone, zone alias or zone set. For example, if you create a zone named "new", you cannot give a zone, zone alias, or zone set the same name.
- The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign (\$).

Naming Conventions for McDATA and Connectrix Switches:

- The name can have a maximum of 64 characters.
- The first character of a zone name must be a letter (A-Z, AZ).
- A zone name cannot contain spaces.
- Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).
- All names must be unique and may not differ by case. For example, myzone and MyZone are considered to be the same zone.


4. Click **OK**.

The new zone is added to the Zone pane.

5. Click the **Create Job** button.



The job is added to the bottom pane.

6. Do one of the following:


- If you want the job to execute now, click the Execute Job () button
- If you want the job to execute at a later time, schedule the job as described in the topic, "[Scheduling Provisioning Jobs](#)" on page 313.

Volume Creation and LUN Security

Keep in mind the following:

- If you have options still selected from a previous job, just clear the options you do not want in your next job. For example, assume you created a volume. Now you want to create a new volume on the same host used previously. Just clear the Volume pane. To clear a pane, click the  button.
- If you want to clear all the steps, except for the Step 1 (storage systems), select another option from the System Action combo-box.
- (HDS only) Before you can create a volume. You must delete some unmapped LDEVs using the standard provisioning tool. To delete LDEVs, click **Provisioning** (). Then, click the Storage Systems tab, then the **Provision** button for the storage system. Then, click **Step 2 Volume**. Click the **Unmapped Volumes** folder in the left pane. Select the desired number of LDEVs for the LUSE volume then click the **Delete Selected Volumes** button. Take note of the array group from which you deleted the LDEVs. You need this information to create the LUSE volume.

You can create a meta volume and designate LUN security, as described in the following steps.

1. Click **Provisioning** ().
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Select the following from the System Action combo-box: Volume Creation + LUN Security

Step 1 - Select Storage System

1. Wait for the management server to load the storage systems into the Storage System panel.
2. Select the storage system on which you want to create the metavolume.

NOTE: The S heading in the Storage Systems pane means that only a single selection is allowed.

(HDS only) Select the storage system from which you want to create the LUSE volume.

3. Click the **Step 1** button underneath the pane.

Information about the selected storage system, such as its model and IP address is displayed below the Storage System pane. The Host pane becomes populated. Notice in the following pane that some hosts have a red X over their icon. This means the host is not accessible.

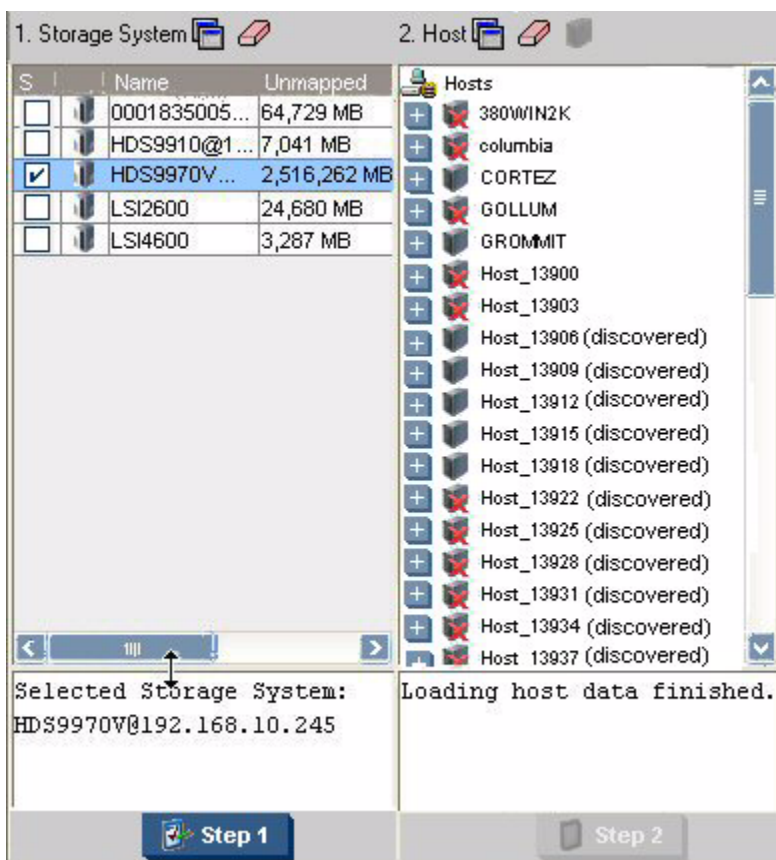


Figure 56 Selecting a Storage System

Step 2 - Select a Volume

To select a volume:



- In the Volume pane select mapped and unmapped volumes. You can select multiple volumes on Windows computers by pressing the Ctrl key as you select the volumes.
 - Mapped** - There are two types of mapped volumes:
 - Masked** - The volume is exposed to the storage port and to the host.
 - Unmasked** - The volume is exposed to the storage port, but not to the host.
 - Unmapped** - The volume is not exposed to the storage port. The management server puts all unmapped volumes in this category when it first discovers an array.
 - Free Extents** - Available free extents that can be used to create a meta volume. You can create meta volumes on EMC Symmetrix and LUSE on HDS storage systems. To create a meta volume or LUSE, select multiple free extents under the Free Extents node in the Volume pane. Select multiple LDEVs from the Free Extents menu by holding down the shift key on

your keyboard and selecting free LDEVS. When you select free extents, they must be of the same type. For example, on Symmetrix, you cannot select a mirrored volume and a BCV (business continuous volume) to create a meta volume.

When you first discover a storage system, no free extents are displayed. This is because the management server puts all unmapped volumes into the “unmapped” category for the list of volumes by default. To move a volume to the free extent node, delete the unmapped volume. When the volume is deleted, it is moved to the free extents node. The free extents category is used internally by the management server. See [“Deleting a Storage Volume”](#) on page 268 for more information.

IMPORTANT: Make sure the free extents you select are not being used. Data on the free extents becomes unusable when a meta volume is created

Keep in mind the following:

- You can narrow the type of volumes displayed in the Volumes pane by using the Volume Customize Dialog () located above the Volumes pane. See [“Host Customize Dialog”](#) on page 317.
- If the LUN has already been selected and Step 4 is clicked, skip this step or click the  button.

2. Click **Step 3**.

3. If you are asked to specify a LUN number, provide a LUN for each volume displayed. See [“Providing a LUN Number”](#) on page 310 for information about numbering LUNs.

Step 3 - Select a Host Security Group


1. Select a host security group in the LUN pane. See [“Creating a Host Security Group”](#) on page 311 for information on how to create a host security group. See [“General Provisioning Issues”](#) on page 282 for information on how your storage system handles host security groups.

2. Click **Step 4**.

3. Click the **Create Job** button.

The job is added to the bottom pane.

4. Do one of the following:

- If you want the job to execute now, click the Execute Job () button
- If you want the job to execute at a later time, schedule the job as described in the topic, [“Scheduling Provisioning Jobs”](#) on page 313.

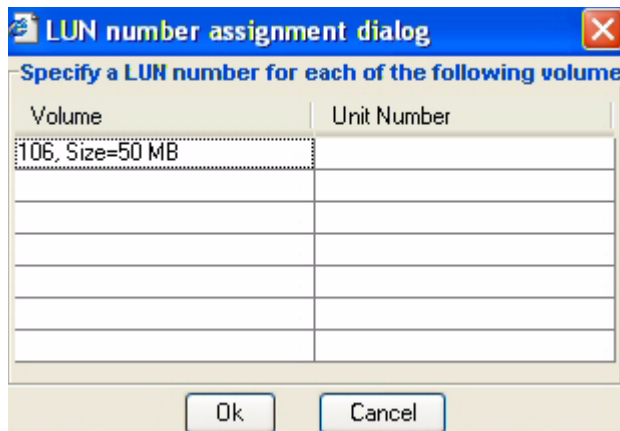
Providing a LUN Number

Keep in mind the following when providing a LUN number:

- **Unique numbers:** LUN numbers must be unique.
- **HDS 9200 series storage systems:** LUN numbers must be between 0 and 128.
- **HDS 9500V series storage systems:** LUN numbers must be between 0 and 512.
- **Engenio storage systems:** LUN numbers must be between 0 and 31.

- **Symmetrix storage systems:** LUN numbers must be between 1 and 8190.

You can enter a LUN number for a volume by placing the cursor under the LUN Number column, as shown in the following figure. Click **OK** when you are done. The window displayed in the following figure is displayed if your storage system requires you to provide a LUN number.



The dialog box is titled "LUN number assignment dialog" with a close button (X) in the top right corner. Below the title bar, it says "Specify a LUN number for each of the following volume". There is a table with two columns: "Volume" and "Unit Number". The first row of the table contains the text "106, Size=50 MB" under the "Volume" column and is empty under the "Unit Number" column. There are four more empty rows below it. At the bottom of the dialog box, there are two buttons: "Ok" and "Cancel".



Volume	Unit Number
106, Size=50 MB	

Figure 57 Specifying a LUN Number

Adding a Host

The management server lets you add hosts that are not currently connected to the network. While you are creating a job, add the host. This feature is only available when you select **LUN Security** from the System Action menu.

IMPORTANT: Make sure the added host is physically connected to the network before the scheduled job runs.


1. Click the  button.
2. Type the name for the host in the **Host Name** field.
3. Type the port WWN of the host in the **Port WWN** field.
4. Click the **Add** button.
5. Repeat Steps 3 and 4 for multiple ports.
6. If you want to remove the host, click the  button.
7. When you are done with your changes, click **OK**.
The host is added to the list of hosts.
8. Physically connect the host to the network before the job runs.

Creating a Host Security Group

You cannot assign the host mode for an IBM storage system in the user interface. You must modify an internal property to set the host mode. See ["Setting the Host Mode for IBM Storage Systems"](#) on page 281 for more information.

IMPORTANT: Each storage system handles host security groups differently. See "[General Provisioning Issues](#)" on page 282.

To create a host security group:

1. Select a storage system in the Storage System pane.
2. Click **Step 1**.
3. Select a host in the Host pane.
4. Click **Step 2**.
5. Select a port in the LUN pane.
6. Click  at the top of the LUN pane.
7. When you are asked to provide a name for the new host security group, type a unique name for the host security group.

NOTE: You are not asked for the name of the host security group for Symmetrix storage systems.

Keep in mind the following:

- The name must contain 1 to 50 characters. If you enter no characters, you are given the option of using a default name.
- The first and last letter cannot be spaces
- You cannot have the following characters in the name:
 - <
 - >
 - ;
 - :
 - ,
 - |
 - /
 - *
 - ?
 - \
 - \\
 - \t
 - \n
 - \b
- 8. (HDS) Select the host mode for the host security group.
- 9. (HDS) Provide a second host mode if applicable.


10.(HDS) Click the **Create Host Security Group** button.

11.For non-HDS storage systems, click **OK**.


The host security group is created in the LUN pane.

Scheduling Provisioning Jobs

Keep in mind the following:

- You must have already created a provisioning job. See "[About Path Provisioning](#)" on page 289 for more information.
- You cannot delete a job once it has started. A job has started if it has a state of "started."
- You can deselect jobs by clicking the check marks or by clicking the Clear Selection () button in the Provision Jobs pane.

You can determine the status of a job by looking in the State column:


- **Created** - The job has been created, but it will not be executed. The job cannot be viewed by others and it is deleted when the Web browser is closed. See "[Scheduling Provisioning Jobs](#)" on page 313 for information about changing the state of the job from "created" to "scheduled".
- **Scheduled** - The job has been tasked to execute at a specified time and date. Jobs are assigned a scheduled state after you select the job and click the **Execute Job** button ( button).
- **Started** - The job has started. You cannot delete a job once it has started.
- **Finished** - The job has finished.

To schedule a provisioning job:


1. Click the **Create Job** button in the lower pane.


The job is given the "created" status. The job, however, is not executed.



Keep in mind the following:

- When you close the Web browser window, all jobs with a status of "created" are erased.
- Other users cannot see a job with a status of "created".
- Change the schedule of job only if its status is "created". Once you click the Execute Job () icon, the job is saved in the database for the management server and its status changes from "created" to "scheduled". Thus, it should not be modified.



If you are unable to click the **Create Job** button, verified all required Step buttons have been clicked. If you are still unable to click the **Create Job** button, verify if the Host Customize

Dialog () is selected for multipathing. If the Host Customize Dialog is set for multipathing, select the second path from the Path combo box and repeat the provisioning steps, except the steps for selecting a system action, storage system and host. See "[Host Customize Dialog](#)" on page 317.

2. Schedule the job by selecting the job and then clicking the  button in the Provision Jobs pane.

- a. (Optional) In the **Time** field, change the time displayed. The management server automatically displays a time five minutes from when you clicked the . Type a time in 24-hour format with hours and minutes separated by a colon. For example, 23:15, for 11:15 p.m.
 - b. Select a date you want the job to start.
The date is selected.
 - c. Click the **Set** button.
If you click the **Set** button after the time has passed, you must reset the time.
3. To execute the job, select the job and then click the Execute Job () button. The job is executed at the scheduled time.
- You can obtain more information about a schedule job by selecting it. When you select a scheduled job in the Provision Jobs pane, detailed information appears in the Job Console pane, located in the lower-right corner. The Message Console tab, located also in the lower-right corner, provides information such as if the job has ended or failed. If a job has failed, the reason for the failure is provided on the Message Console tab.
- Keep in mind the following:
- Jobs are executed according to the time set on the management server, since they are centrally saved in the database for the management server.
 - When the management server is determining which job to perform first, it looks first for jobs requiring volume and/or zone creation. It does not perform the jobs in the order suggested by the Provision Job ID nor according to the order of the jobs appearing in the Provision Jobs table.
 - The management server can detect when the required volume and/or zone has already been created. For example, assume you created a job to create a volume. The next few jobs use this volume. These jobs will determine if the volume has been created, before creating the volume.

Executing Provisioning Jobs


If you want to save and execute a job, you must click the **Execute Job** button () button. When you click the **Execute Job** button () button, the job is saved on the management server. Other users can now see the job.

When the management server is determining which job to perform first, it looks first for jobs requiring volume and/or zone creation. It does not perform the jobs in the order suggested by the Provision Job ID nor according to the order of the jobs appearing in the Provision Jobs table.

The management server can detect when the required volume and/or zone has already been created. For example, assume you created a job to create a volume. The next few jobs use this volume. These jobs will determine if the volume has been created, before creating the volume.

You can schedule a job to start now or at a pre-determined time in the future:


- **To start a job now** - Click the **Execute Job** button () button before the job is scheduled.

- **To start a job at a pre-determined time in the future** - Schedule the job first. Then, click the **Execute Job** button () button. The job is executed according to the time set on the management server.


To view the latest information in Path Provisioning, click the **Refresh** button. The **Refresh** button updates the following:


- The Path Provisioning screen with changes made in Provisioning Manager. For example, assume you used the wizards in Provisioning Manager to create a host security group. When you access Path Provisioning, your changes are not shown.
- The Path Provisioning screen with changes from executed jobs. After a job is executed in Path Provisioning, the Path Provisioning screen is not updated until you click the **Refresh** button or you exit and re-enter Path Provisioning.
- Other parts of the product, such as:
 - Application Viewer
 - Capacity Manager
 - Performance Manager
 - Provisioning Manager
 - System Manager

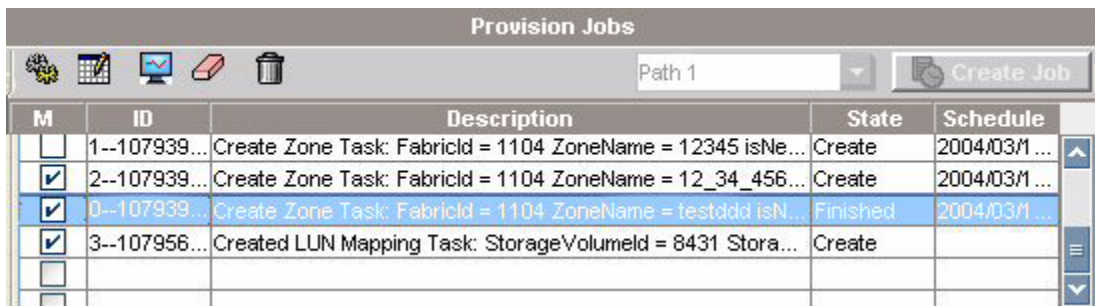
Monitoring Provisioning Jobs

If you want to view the latest status of the provisioning jobs listed, click the  button in the Provision Jobs pane. The management server gathers information about the provisioning jobs listed to determine their latest status. You might want to try this feature when you are not sure if a job has finished.

Deleting Multiple Jobs

IMPORTANT: You cannot delete a job once it has started. A job has started if it has a state of “started.” To delete jobs, select the jobs in the M column and then click the  icon until all jobs have been removed, as shown in the following figure.

You can deselect jobs by clicking the check marks or by clicking the Clear Selection () button in the Provision Jobs pane.



M	ID	Description	State	Schedule
<input type="checkbox"/>	1--107939...	Create Zone Task: FabricId = 1104 ZoneName = 12345 isNe...	Create	2004/03/1 ...
<input checked="" type="checkbox"/>	2--107939...	Create Zone Task: FabricId = 1104 ZoneName = 12_34_456...	Create	2004/03/1 ...
<input checked="" type="checkbox"/>	0--107939...	Create Zone Task: FabricId = 1104 ZoneName = testddd isN...	Finished	2004/03/1 ...
<input checked="" type="checkbox"/>	3--107956...	Created LUN Mapping Task: StorageVolumeld = 8431 Stora...	Create	
<input type="checkbox"/>				
<input type="checkbox"/>				

Figure 58 Deleting Multiple Provisioning Jobs

Naming Conventions for Zones

The following are the naming conventions for zones:

Naming Conventions for Brocade Switches:


- The name must contain 1 to 64 characters.
- The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).
- The name is case sensitive. For example, “Zone1” and “zone1” are different zones.
- You cannot create a zone with the same name as an existing zone, zone alias or zone set. For example, if you create a zone named “new”, you cannot give a zone, zone alias, or zone set the same name.
- The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign (\$).

Naming Conventions for McDATA and Connectrix Switches:

- The name can have a maximum of 64 characters.
- The first character of a zone name must be a letter (A-Z, AZ).
- A zone name cannot contain spaces.
- Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).
- All names must be unique and may not differ by case. For example, myzone and MyZone are considered to be the same zone.

Using Multipathing with Path Provisioning

You can use provisioning with multipathing. Set the Host Customize Dialog to the multipathing option. Then, repeat the provisioning steps for each path, as described in the following steps:

1. Select one of the following system actions:
 - **Volume Creation + LUN Security + Zoning** - First you create a meta volume. Second, you map a meta volume to a Fibre Channel port and host HBAs (HSG). Third, you create a zone. See "[Volume Creation, LUN Security, and Zone Operation](#)" on page 293 for more information.
 - **LUN Security** - Map meta volume to Fibre Channel port and host HBAs (HSG). See the topic "[LUN Security](#)" on page 299 for more information.
 - **Zone Operation** - Create a new zone. See "[Zone Operation](#)" on page 304 for more information.
2. Select a storage system, and then click the **Step 1** button.
3. Click the  button above the Host pane.
4. Select the following option:
`Multipath: Select more than one port within a single server.`
5. Select a port on a host. Then, click the **Step 2** button.
6. Select a volume, host security group and/or zone, as described in the following topics:
 - "[LUN Security](#)" on page 299
 - "[Zone Operation](#)" on page 304
7. Select the second path from the Path combo box, as shown in the following figure:
8. Repeat Step 6.
You do not need to select a system action, storage system, and host. Furthermore, you must select the same host used for the first path.
9. Set the schedule for the job as described in the topic, "[Scheduling Provisioning Jobs](#)" on page 313.

You can narrow the types of volumes displayed in the Volume pane by setting the Volume Customize Dialog.

Customizing Path Provisioning

This section describes the following:

- "[Storage System Customize Dialog](#)" on page 317
- "[Host Customize Dialog](#)" on page 317
- "[Volume Customize Dialog](#)" on page 318
- "[LUN Customize Dialog](#)" on page 318
- "[Zone Customize Dialog](#)" on page 318

Storage System Customize Dialog

The Storage System Customize Dialog lets you specify the type of storage systems displayed in the Storage System pane. Select one of the following:

- **Show all available storage systems** - All storage systems are displayed.
- **Show storage system with available raw storage** - Only storage systems with available raw storage are displayed.
- **Show storage system with available unmapped storage** - Only storage systems with available unmapped storage are displayed.
- **Show storage system with the following characteristics** - Select one or more of the storage tiers. Only storage systems belonging to the selected storage tiers are displayed. Storage systems are assigned to storage tiers in Chargeback Manager. See "[Adding Asset Information](#)" on page 463 and "[Adding General Information](#)" on page 464 for information about assigning storage systems to a storage tier.

Click **OK** when you are done with your selections.

The Storage System pane is updated.

Host Customize Dialog

The Host Customize Dialog lets you use multipathing with Path Provisioning. Select one of the following options:

- **Single path: Select one port of a host.** Select this option if you do not have multipathing or you do not want to use multipathing with Path Provisioning.
- **Multipath: Select more than one port within a single server.** Select this option if you want to use multipathing with Path Provisioning. See "[Using Multipathing with Path Provisioning](#)" on page 316 for more information on how to use multipathing with Path Provisioning.

Volume Customize Dialog

1. If you want to view all volumes, select the **Show All Volumes** option. Once you have selected the **Show All Volumes** option, you have the option to select the **Hide mapped volumes** option, which hides mapped volumes.
2. Select one of the following for metavolumes:
 - **Concatenating** - Only concatenating metavolumes are displayed.
 - **Strip** - Only stripped metavolumes are displayed. (Applies only to EMC storage systems)
3. Click **OK** when you are done with your selections.

The Volume pane is updated.

LUN Customize Dialog

1. Specify how LUN are mapped by selecting one or more of the following options in the LUN customize dialog:
 - **Perform number of path verification based on host selection** - The path verification is based on the host you selected.
 - **Automatically Assign Volume to Storage Port based on** - Select one of the following:


- ***Most unused ports (the number of LUNs assigned)** - Assigns a volume to a FA port based on how often a port is used. Unused ports have a better chance of having a volume assigned to them as compared to frequently used ports. This option helps you spread out the traffic.
- ***Linked port if there is any** - A linked port is more likely to be assigned a volume. The management server looks for any storage system port that is zoned to the selected host bus adapter. If the management server cannot find a storage system port zoned to the selected HBA, the management server selects a port with the least amount of LUNs.

*You can only access these options if **Automatically Assign Volume to Storage Port Based on** is selected.

- **Verify that in a multipath configuration, storage system ports do not connect to the same switch** - This option makes sure the multipath configuration is preserved. A storage system ports should not get connected to the same switch.
- **Assign a LUN number automatically** - Do not select this option if you want to assign a LUN number manually.

2. Click **OK** when you are done with your selections.

Zone Customize Dialog

When the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**, is selected in the zone customize dialog () , the management server automatically selects a zone that meets its criteria. If the management server cannot find a zone that meets its criteria, it creates a zone on the fly. You set the criteria for automatically configuring zones by selecting one of the following options:

- ***Create new zone if there is no zone containing ONLY the selected zone members (HBA and storage system ports)** - The management server checks if an identical zone exists. An identical zone contains only same HBA and storage system ports you selected. If the zone contains additional members, it is not considered to be identical.
 - If an identical zone exists, it is selected in the Zone pane.
 - If an identical zone does not exist, the management server asks you to name a zone that will be created. The new zone appears in the Zone pane, even though it is not created until the job runs.
- ***Create new zone if there is no zone containing AT LEAST the selected zone members** - The management server tries to find a zone that contains the HBA and storage system ports you selected.
 - If a zone contains additional members, the management server selects that zone in the Zone pane.
 - If the management cannot find a zone containing the HBA and storage system ports you selected, the management server asks you to name a zone that will be created. The new zone appears in the Zone pane, even though it is not created until the job runs.

*To enable one of these options, the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**, must be selected.

IMPORTANT: (McDATA switches only) Path Provisioning looks for the names of the active zone set and of the active zones and all of their saved counterparts in the zoning library in EFC Manager. The provisioning job only occurs if those names match.


About the Use Switch Port Zoning Mode Option

By default Path Provisioning creates zones through Worldwide Names, which means it looks for a port on a host and a port on a storage system to create the zone. Zoning through WWNs is not dependent upon the switch. This means you could change switches and as long as the host and storage system are able to access each other through the network.

You can create zones through switch port zoning if you select the **Use Switch Port Zoning Mode** option. When you select the **Use Switch Port Zoning Mode** option, you select ports on a switch to create your zone. For example, ports 1 and 2 on a switch could be designated for a zone. This means that whatever is connected to either of those ports will be a part of the zone.

Automatically Configure Zoning

To automatically configure zoning:

1. Click the  button above the Zone pane before you select a host.
2. Select the option, **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**.
3. Select one of the options mentioned in the previous bulleted list. Then, click **Apply**.
4. Select a storage system and click **Step 1**.
5. Select a host and click **Step 2**.
6. One of the following occurs:
 - If the management server cannot find a zone that meets the criteria set in Step 3, it asks for a zone name.
 - If the management server finds a zone that meets the criteria set in Step 3, it selects the zone in the Zone pane.
7. If you are asked for a zone name, type the zone name and then click **OK**.


See "[Naming Conventions for Zones](#)" on page 316 for restrictions on naming zones.

The new zone is displayed in the Zone pane, but it has not been created yet. It will be created when the job runs.

Manually Configure Zoning

The management server assumes you want to select a pre-existing zone or create one manually when you deselect the **Automatically create new zone if no existing zone containing HBA and storage system ports is detected** option.

- **To reuse a zone** - Select a zone in the Zone pane and then click **Step 5**. Expand a fabric node to view its zones.

- **To create a zone** - Select a fabric in the zone pane, click the  button, type a zone name. See the following topic for more information.






If you see the  icon, the fabric cannot be reached.

Table 51 Zone Icons

Icon	Description
	Its location on the page determines its use: <ul style="list-style-type: none"> • Displayed above Zone pane - Button for creating zones. • In Zone pane - Icon for a zone.
	Zone Alias
	Port

To create a zone:

1. Select a fabric in the zone pane.
2. Click the  button located above the Zone pane.
3. Type a zone name in the dialog box. The following restrictions apply:

Naming Conventions for Brocade Switches:

- The name must contain 1 to 64 characters.
- The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).
- The name is case sensitive. For example, "Zone1" and "zone1" are different zones.
- You cannot create a zone with the same name as an existing zone, zone alias or zone set. For example, if you create a zone named "new", you cannot give a zone, zone alias, or zone set the same name.
- The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign (\$).

Naming Conventions for McDATA and Connectrix Switches:

- The name can have a maximum of 64 characters.
- The first character of a zone name must be a letter (A-Z, AZ).
- A zone name cannot contain spaces.
- Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign (\$).
- All names must be unique and may not differ by case. For example, myzone and MyZone are considered to be the same zone.

4. Click **OK**.

The new zone is added to the Zone pane.

11 Event Management

This chapter describes the following:

- ["About Event Monitoring for Storage Essentials"](#) on page 323
- ["Viewing Event Details"](#) on page 326
- ["Clearing Events"](#) on page 327
- ["Configuring the Clearing of Events"](#) on page 327
- ["Configuring the Deletion of Events"](#) on page 328
- ["Deleting Events"](#) on page 329
- ["Sorting Events"](#) on page 329
- ["Changing the CLARiiON Event Polling Interval"](#) on page 330
- ["Brocade Events"](#) on page 330
- ["Filtering Events"](#) on page 332
- ["Enabling the Rank Column in Event Monitoring for Storage Essentials"](#) on page 335

About Event Monitoring for Storage Essentials

Event Monitoring for Storage Essentials lets you view, clear, sort and filter events from managed elements. An event can be anything that occurs on the element, such as a device connected to a Brocade switch has gone off-line. It provides the following information about the events:

- **ID** - The identification number assigned to the event
- **Severity** - Provides the severity level
- **Time** - The time the event was recorded.
- **Summary Text** - A brief explanation of the event. When you click the summary text, the details of the event are displayed.

To view events for a specific element:

- Double-click the element in Capacity Manager, Performance Manager or System Manager and then click the **Events** tab. Only events from the element that was double-clicked are displayed. You cannot access the **Element Type** filter through this method.
- Right-click the element in Capacity Manager, Performance Manager or System Manager and then select the **Show Events** option in the menu.
- In HP Systems Insight Manager, click the **Storage Essentials Managed** link. Click the link for the system name. On the HP System Page, click **SE System Properties**. Then, click the **Events** tab.








You can change your user preferences for Event Monitoring for Storage Essentials by doing one of the following:

- Clicking the Preferences link in Event Monitoring for Storage Essentials. See ["Modifying Your User Preferences"](#) on page 81 for more information.

- Clicking **Options > Events > Storage Essentials > Event Configuration**. See ["Controlling the Display of Cleared and Deleted Events"](#) on page 110.

The following icons are displayed in Event Monitoring for Storage Essentials.

Table 52 About the Icons in Event Monitoring for Storage Essentials

Icon	Description
	Event was marked cleared. See the topic, "Clearing Events" on page 327 for more information.
	The severity of the event is not known.
	The event is informational.
	The event might have some impact.
	The event has a minor impact.
	The event has a major impact.
	The event has a critical impact.

Event Monitoring for Storage Essentials does not support events from all discovered elements. See the following table for more information about which elements Event Monitoring for Storage Essentials supports.

Table 53 Supported Hardware

Hardware	Events Supported?
Brocade switches	Y, 1
4QLogic SNMP switches	Y, 3
5McDATA switches SWAPI to EFCM	Y,2
5McDATA switches SNMP through proxy	Y, 3
5McDATA switches SNMP to switches	Y, 3
Cisco switches SNMP	Y, 3
CNT switches SMI	N
Engenio and Sun 6130 storage systems	Y
CLARiiON storage systems	Y
Symmetrix storage systems	Y

Table 53 Supported Hardware (continued)

Hardware	Events Supported?
HDS storage systems	Y
Sun 35xx storage systems	N
Sun 6920 storage systems	N
HP-XP SMI storage systems	N
HP-MSA storage systems	N
HP-EVA storage systems	N
IBM ESS storage systems	N

Notes from the table:

1Not all events that show up in the webtool appears in Event Monitoring for Storage Essentials.

2Not all events that show up in EFCM will appear in Event Monitoring for Storage Essentials.

3Need to configure the switch or proxy to send traps to the management server.

4Also applies to Sun StorEdge switches.

5Also applies to EMC Connectrix switches.

Keep in mind the following:

- The management server provides separate event notification for every event that is reported from the devices it is monitoring. Excessive notification could delay provisioning, as the providers are busy notifying the management server of the events. If you do not want the management server to be notified of every event, modify the event threshold of the devices to filter out some of the events. Refer to the documentation accompanying the device for more information about setting the threshold.
- To receive events from Sun StorEdge or QLogic switches, verify the SNMP trap community string is set to public in SANbox Manager or through telnet. Also, make sure the SNMP traps are configured to be sent to the management server.
- The management server does not receive SNMP v1 traps from Sun StorEdge switches that have the SNMPv1 agent.
- If you want the management server to be able to receive events from a NetApp Filer, you must add the IP address of the management server CIMOM to the NetApp configuration. The management server CIMOM runs on the same computer running the management server by default.
- Event Monitoring for Storage Essentials displays events it receives from CNT InVsn Enterprise Manager. As of version 9.5 of CNT InVsn Enterprise Manager, InVsn Enterprise Manager does not provide events to the management server. As future versions of CNT InVsn Enterprise Manager provide event support, the management server will be able to provide information about those events.

Issues with Brocade Switches

- Event Monitoring for Storage Essentials does not display events from Brocade switches with the firmware version 3.0. This firmware version is not supported by Event Monitoring for Storage Essentials. You can, however, specify Event Monitoring for Storage Essentials not to display events from additional firmware versions. See ["Filtering Events"](#) on page 332.
- Events from Brocade Fabric Watch have "Fabric watch:" in their heading, as shown in the following example:

```
Fabric watch: threshold is below low boundary, Value:30 Class: Environment
ClassArea: Temperature Index: 1 EventType: Below EventState: Faulty
```

Issues with McDATA and Connectrix Switches

- If you discovered McDATA and/or Connectrix switches through SWAPI, Event Monitoring for Storage Essentials does not report events for switch hardware failures, except for those regarding switch ports (port off-line/port online).
- If you are not receiving events from McDATA and Connectrix switches discovered through SNMP, make sure you have the correct port set for receiving SNMP traps. You must also configure the proxy or each switch (depending on your configuration for discovery) to send the SNMP traps to the correct port on the management server. See ["Discovering McDATA and EMC Connectrix Switches"](#) on page 18 for more information.
- If you want the management server to receive SNMP events from Connectrix or McDATA switches, do one of the following:
 - If you discovered Connectrix Manager or EFC Manager, only enable SNMP trap forwarding to the management server on the Connectrix Manager or EFC Manager, not on the individual switches. Connectrix Manager or EFC Manager should be configured to forward SNMP traps to the IP address of the management server, and the community string should match the user ID you used to discover Connectrix Manager or EFC Manager.
 - If you discovered Connectrix or McDATA switches directly, enable SNMP trap forwarding on the switches, not on any other management software. The switches should be configured to forward SNMP traps to the IP address of the management server, and the community string should match the user ID you used to discover the Connectrix or McDATA switches.

Viewing Event Details

To access event details:

1. Access Event Monitoring for Storage Essentials as described in ["About Event Monitoring for Storage Essentials"](#) on page 323.
2. Click the summary text for the event.

The event details are displayed.

The Event Details pane provides information on one or more of the following:

- **Serial Number** - The number assigned to the event
- **Summary Text** - A brief explanation of the event
- **Element** - It is the source of the event. An element can be a switch, host, application, fabric or anything else on the network. If this field is blank, the event did not come from an element.
- **Source** - If the event came from an element, the name of the element is provided.

- **Time Reported** - The time the event was reported to the management server.
- **Probable Cause** - Provides an explanation of a probable cause.
- **Recommended Actions** - Provides recommendations.
- **Event Type** - Provides a description of the event.
- **Correlated Indications** - Information obtained from the provider.
- **Severity** - Provides the severity level, which can be one of the following:
 - **Clear**
 - **Unknown**
 - **Informational**
 - **Warning**
 - **Minor**
 - **Major**
 - **Critical**

NOTE: Events listed in Event Monitoring for Storage Essentials may not be attributed to the correct source until Discovery Data Collection has completed.

Clearing Events

You might want to mark an event as “cleared” if you have already reviewed it. Use one of the following techniques to mark an event as “cleared”:

- Click the **Clear this Event** button in the Event Details pane, which is accessible by clicking the summary text link.
- Right-click the event listed in Event Monitoring for Storage Essentials and select **Mark as Cleared** from the drop-down menu. This menu does not appear if you right-click the summary text link.
- Select the check box to the left of the event and click the **Clear** button in Event Monitoring for Storage Essentials.

Note: You can clear more than one event by selecting the check box to the left of the ID column header and then clicking the **Clear** button in Event Monitoring for Storage Essentials.

Unclearing Events

You can unclear events by doing one of the following:

- Click the **Set this Event to Not Cleared** button in the Event Details pane
- Select the check box to the left of the event and click the **Undo Clear** button in Event Monitoring for Storage Essentials. You can unclear multiple events at once by select them and then clicking the **Undo Clear** button.
- Right-click an event in Event Monitoring for Storage Essentials and select the **Mark as NOT Cleared** option from the drop-down menu. This menu does not appear if you right-click the summary text link.

Configuring the Clearing of Events

Depending on the severity of an event, the management server may mark the event as “clear” after 60 minutes. Events designated as Major and Critical are never marked as “clear.” You can change the time delay in clearing an event, and you can specify that the management server never marks an event as “clear”.

To help you in filtering events, you may want to have unimportant events marked as cleared rather than automatically deleted. Depending on how you have configured the deletion of events, you can view the “cleared” events at a later time.

See the following table for the default settings for clearing events.

Table 54 Default Settings for Clearing Events

Severity Level	Default Time Delay to Clear the Event (Hours)
Unknown	1
Informational	1
Warning	1
Minor	1
Major	Never
Critical	Never

To change the default time delay to clear an event, do the following:

1. Select **Options > Events > Storage Essentials > Event Configuration**.
2. Do one of the following:
 - If you never want an event of the specified severity level marked as “cleared,” select the **Never** option next to the severity level.
 - If you want to change the delay time in clearing an event, select one of the following units of measurement from the combo box and type the number in the adjacent field:
 - **Minutes**
 - **Hours**
 - **Days**
 - **Weeks**
3. Click the **Save Changes** button.

Configuring the Deletion of Events


The management server automatically deletes events after two weeks by default. You can specify for each severity level different time periods for deleting events. For example, you could modify the management server to delete events with the Information severity level every two days. You could also specify the management server to never delete events with the Critical severity level.

To change the default time delay to delete an event, do the following:

1. Select **Options > Events > Storage Essentials > Event Configuration**.
 2. Do one of the following:
 - If you never want an event of the specified severity level automatically deleted, select the **Never** option under the Automatic Delete Delay column.
 - If you want to change the delay time in deleting an event, select one of the following units of measurement from the combo box and type the number in the adjacent field:
 - **Minutes**
 - **Hours**
 - **Days**
 - **Weeks**
- For example, if you want events that are a week old deleted, you would type 1 and then select **Weeks** in the combo box in the Automatic Delete Delay column.
3. Click the **Save Changes** button.

Deleting Events

To delete an event:

1. (Optional) Select the events you want to delete in Event Monitoring for Storage Essentials.
2. Click the **Delete** button at the top of the screen.
3. Select one of the following options in the Select Events to Delete window:
 - **Delete selected events** - All selected elements are deleted.
 - **Delete all events** - All elements are deleted, regardless of your filter settings.
 - **Delete events earlier than** - Elements earlier than the date specified are deleted, regardless of your filter settings. Click the calendar icon, .
4. Click **OK**.

Sorting Events

In Event Monitoring for Storage Essentials, you can sort events. For example, if you want to see the most severe event on a page, click the Severity column header link. Click it again to sort events in the reverse order.

To sort the events:

1. Access Event Monitoring for Storage Essentials as described in "[About Event Monitoring for Storage Essentials](#)" on page 323.
2. In the Event Monitoring for Storage Essentials table, click the column title corresponding to the attribute you want to sort the event. The following is a listing of the attributes:
 - **ID** - The identification number assigned to the event
 - **Element** - The source of the event. An element can be a switch, host, application, fabric or anything else on the network.

- **Severity** - Provides the severity level
- **Time** - The time the event was recorded.
- **Summary Text** - A brief explanation of the event. When you click the summary text, the details of the event are displayed.
- **Element Type** - Specifies whether the source of this event is an application, a host, etc.
- **Rank** - The cost implication. This column is hidden until you enable it, as described in ["Enabling the Rank Column in Event Monitoring for Storage Essentials"](#) on page 335.

Changing the CLARiiON Event Polling Interval

You can change how frequently the management server polls the CLARiiON storage systems by modifying the `cimom.ClariionEventPollInterval` property. You may want to change this interval if you are receiving many "information" messages from the CLARiiON storage system.

IMPORTANT: Do not set a very long time interval. The management server does not become aware of events occurring on CLARiiON storage system until it polls the storage system. For example, assume you set the polling interval to every two days. A serious issue could occur on the first day, but you would not know about it until the second day because you set a long time interval.

To change the polling interval:

1. Click **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the `cimom.ClariionEventPollInterval` property. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
4. Return to the Advanced page (**Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make your changes in the **Custom Properties** field. Make sure the property is not commented out by removing the hash (#) symbol in front of the property.
7. Change the value assigned to the `cimom.ClariionEventPollInterval` property. The value is in milliseconds. In the example below, the polling interval is set to 5 minutes.

```
cimom.ClariionEventPollInterval=300000
```
8. When you are done, click **Save**.
9. Restart the service for the management server for your changes to take effect.

While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Important: While the AppStorManager service is stopped, the following occurs:

- Users are not be able to access the management server.
- The management server is unable to monitor elements at this time.

Brocade Events

This section describes the following:

- “[Brocade Switch Events](#)” on page 331
- “[Supported Brocade Events](#)” on page 332

Brocade Switch Events

When a Brocade switch generates an event, it assigns a code to the event instead of an event severity level. So, the software assigns an event severity level to the event according to the event's code. This lets you filter Brocade switch events by severity level in Event Monitoring for Storage Essentials, as described in the following table.

NOTE: Events regarding firmware downloads are removed from the following table since the management server cannot be made aware of those events.

Table 55 Brocade Switch Events

*Code	Event Severity Level	*Name	Description
0	Minor	EV_OBJ_CHANGED	A physical fabric object (switch port or fabric) has changed state.
1	Major	EV_OBJ_DELETE	One or more physical fabric objects (device port, switch, or fabric) have disappeared.
2	Warning	EV_OBJ_CREATE	One or more new physical fabric objects (device port, switch, or fabric) have appeared.
3	Critical	EV_CONNECTED_OBJECT_OFFLINE	A device connected to a switch has gone offline.
4	Major	EV_CONNECTED_OBJECT_ONLINE	A device connected to a switch has come online.

Table 55 Brocade Switch Events (continued)

*Code	Event Severity Level	*Name	Description
5	Info	EV_RSCN	For those RSCN events not covered by EV_OBJ_xxx codes listed above. Examples include: fabric, domain, connected area state unknown, connected device state unknown.
7	Minor	EV_API_HEART_CONDITION	Monitoring proxy switch.
11	Major	EV_STATE_CHANGE	State changes such as: login, logout, login failed, configuration change, track on/off, port up/down, fabric segmentation, security violations, zone change.
12	Major	EV_PLATFORM_CHANGE	The platform database has changed.

*This information does not appear in the event description, but it appears here for additional information.

Supported Brocade Events

The Event Monitoring for Storage Essentials displays the following events from Brocade switches:

- **RSCN events:** Events about the state of the switch, such as it being offline.
- **Fabric Watch events:** Events about the threshold of the switch, such as threshold below.
- **Fabric Access library events:** Events about proxy switch health.
- **Zoning events:** Events about zoning, such as zone-related state change notification.

Filtering Events

This section describes the following:

- ["About Filtering Events"](#) on page 332
- ["Selecting a Severity Level"](#) on page 333
- ["Customizing the Severity Level Filter"](#) on page 334
- ["Filtering Events by Summary Text"](#) on page 334
- ["Hiding Filters in Event Monitoring for Storage Essentials"](#) on page 334

About Filtering Events

The management server provides several types of event filters to specify which events you want Event Monitoring for Storage Essentials to display.

You can use all the event filters at once or you can use just one of them. You can filter events by:

- **Severity level**
- **Summary text**

Once you set all of your filters, you would click the **Apply Filters** button.

Selecting a Severity Level

You can filter the events, so you only view events that are of a specified severity level. For example, you can specify that Event Monitoring for Storage Essentials displays only events with a Critical severity level.

To select a severity level, select an option from the **Show Severity** drop-down menu in Event Monitoring for Storage Essentials. Then, click the **Apply Filters** button. The definition for each severity level varies according to the type of element.

IMPORTANT: Events marked as “Clear” are not displayed if you select a severity level other than **Clear** or **All**.

Table 56 Severity Levels

Severity Level	Description
Clear	The event has been marked as cleared.
Unknown	The event does not fall into the other categories and further information can not be obtained from it.
Informational	Provides informational data. For example, for Brocade switches it could be a list of switches that have successfully completed firmware download.
Warning	Provides warning data. For example, for a Brocade switch one or more new physical fabric objects (device port, switch, or fabric) have appeared.
Minor	Provides a message to indicate a minor problem. For example, for a Brocade switch a physical fabric object (switch port or fabric) has changed state.

Table 56 Severity Levels (continued)

Severity Level	Description
Major	Provides a message to indicate a major problem. For example, for a Brocade switch one or more physical fabric objects (device port, switch, or fabric) have disappeared.
Critical	Provides a message to indicate a major problem. For example, for a Brocade switch a device connected to the switch has gone off line.
All	Displays all events, regardless of severity level.
All But Clear	Displays all events, except for those marked cleared.

Customizing the Severity Level Filter

You can customize the severity level in your filter by clicking the **Customize** button next to the **Show Severity** drop-down menu in Event Monitoring for Storage Essentials.

For example, you can specify you want Critical and Major events displayed in Event Monitoring for Storage Essentials.

IMPORTANT: Events marked as “Clear” are not displayed if you select a severity level other than **Clear**.

To select more than one severity level for filtering:

1. Click the **Custom** button next to the **Show Severity** drop-down menu in Event Monitoring for Storage Essentials.
2. In the Event Filters window, make sure the severity levels you want to appear in Event Monitoring for Storage Essentials are selected.
Notice in the previous figure that the number of events meeting a criteria, such as Minor, is displayed next to the options.
3. Click **OK**.
4. Click the **Apply Filters** button.

Event Monitoring for Storage Essentials displays the events that meet the criteria selected in the Event Filters pane.

Filtering Events by Summary Text

You can filter events by summary text.

To filter elements by summary text:

1. In the **Summary Text** field in Event Monitoring for Storage Essentials, type the summary text or element name you want to be used in the filter.
2. Click the **Apply Filters** button.
Only events with the specified text are displayed in Event Monitoring for Storage Essentials.

Hiding Filters in Event Monitoring for Storage Essentials

Hide the filters for additional screen space. When you hide the filters, the following features are hidden:

- Show Severity
- Summary Text

To hide the filters, click the **Hide Filters** link in the upper-left corner of Event Monitoring for Storage Essentials.

To display the filters, click the **Show Filters** link in the upper-left corner of Event Monitoring for Storage Essentials.

Enabling the Rank Column in Event Monitoring for Storage Essentials

By default, the Rank column in Event Monitoring for Storage Essentials is hidden. The Rank column displays the cost implication of an event, which is determined by the business cost of the application that uses the element and the severity of the event. The business cost of an application is set in System Manager, as described in "[Assigning a Business Cost to an Application](#)" on page 192.

To enable the Rank column:

1. Click **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the `calculateBusinessCost` property. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
4. Return to the Advanced page (**Options > Storage Essentials > Manage Product Health**). Then, click **Advanced** in the Disk Space tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.

6. Change the value assigned to the `calculateBusinessCost` property to `true` in the **Custom Properties** field.
`calculateBusinessCost=true`
7. When you are done, click **Save**.
8. Restart the service for the management server for your changes to take effect.

12 Running Reports

IMPORTANT: Depending on your license, Reporting may not be available. See the “List of Features” to determine if you have access to Reporting. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).

This chapter describes the following:

- “[About Reporting](#)” on page 337
- “[Accessing Reporting](#)” on page 339
- “[Viewing Reports](#)” on page 339
- “[Viewing Report Collectors for an Element](#)” on page 340
- “[Refreshing a Report](#)” on page 341
- “[Changing the Formatting of a Report](#)” on page 341
- “[Opening a Report in a New Window](#)” on page 341
- “[Maximizing the Screen Space for a Report](#)” on page 342
- “[Filtering Data in Global Reports](#)” on page 342
- “[Sending a Report by E-mail](#)” on page 343
- “[Managing E-mail Schedules for Reports](#)” on page 344
- “[Creating Custom Reports](#)” on page 349

About Reporting


Reporting provides a variety of detailed reports, such as dependency, event, and utilization reports for discovered elements. To view a report, click a report name in the tree in Reporting. The report appears in the right pane.

IMPORTANT: If you are shown the message “The report does not contain any data,” verify the collector for the report is running. You may also want to verify a collector is running if you believe the report is not displaying the latest information. Data collectors gather information for reports. If you stop a data collector, its reports are not updated. See “[Managing Collectors for Reports](#)” on page 121.

The management server provides reports that display performance information in a variety of formats:

- **HTML (Default)** - The software displays the report in a Web page by default.
- **PDF** - The software displays the report in Adobe Acrobat, a good option if you need to print the report. The software assumes you already have Adobe Acrobat Reader installed on your computer. To obtain a copy of Adobe Acrobat Reader, go to <http://www.adobe.com>.

- **Excel** - The software displays the report in Microsoft Excel, providing you have a copy of Microsoft Excel already installed.
- **XML** - The software displays the report in the XML format.

It also provides some reports with pie charts. The  icon next to the report name indicates the report displays a graphic.

The software provides reports for the following:

- **Global** - These global-wide reports provides data gathered from multiple management servers. For example, lets assume you have three management servers: one in London, one in Tokyo, and one in New York City. You can gather data from all three management servers. To view information for these reports, you must set up global reporting. To learn more, see ["Setting Up Global Reporter"](#) on page 130 for more information.
- **Asset Management** - These reports provide information based on assets and ownership.
- **Chargeback Manager** - These reports provide cost information about the management and storage usage of an element. To populate these reports, enter information for Chargeback Manager, as described in the topic, ["Setting Up Chargeback Manager"](#) on page 458.
- **System** - These reports are enterprise wide and they collect information about the following:
 - **Application** - Data about applications the management server monitors, such as reports on application utilization and dependencies.
 - **Events** - Data about events occurring on the elements the management server monitors, such as summary reports on events.
 - **Fabric** - Data about fabrics, such as SAN components not zoned and world wide names that appear in zones but not in SANs.
 - **File System Viewer** - Data about the file servers the management server monitors, such as reports on groups and users by server. This information is provided only if you have purchased the license for File System Viewer.
 - **HBA** - A summary report on the host bust adapters (HBAs) the management server detects.
 - **Host** - Data about the hosts in the management server monitors, such as reports on host storage allocation and total host utilization.
 - **Storage System** - Data about storage systems the management server monitors, such as reports on storage system capacity and storage system utilization.
 - **Switch** - Data about switches the management server monitors, such as reports on switch port traffic and port utilization by connection type.
- **Applications** - These reports provide information about an application, such as Oracle or Microsoft Exchange.
- **Hosts** - These reports provide information about a host.
- **Storage System** - These reports provide information about a storage system.
- **Switches** - These reports provide information about a switch.
- **Tape Libraries** - These reports provide information about a tape library.
- **Recent** - Lists the last 10 reports viewed. This option is not displayed when you first access Reporting.

Troubleshooting Reporting

- Reporting does not provide reports for NetApp filers.
- Reports in general contain no data or incomplete/incorrect data while the report cache is being refreshed. See ["Refreshing a Report"](#) on page 341.
- Some reports display data trends for several days in the future. For example, assume you run a report to gather information from the last three days. The report may display data trends for the next several days based on the current information. Please keep in mind the data trends are just assumptions and should not be treated as fact.
- The following IBM HBA appears as QLogic HBAs in the Navigation and Properties pages, in addition to reports:
 - IBM MSJ
 - FaStT FC-2/2-133
- Certain reports display elements assigned to the user's organization, including child organizations. For example, if you attempt to view an Assets by Department report and you do not have permission to access hosts through your organization, you are not given information about the hosts in the report. This is also true for e-mailing certain reports. Let's assume again you do not have permission to access hosts. If you e-mail an Assets by Department report, your e-mail will not contain information about hosts. If the users receiving your reports want to be able to view information about hosts, one of the following must happen:
 - The hosts in question must be added to your organization.
 - Someone else, who has the hosts in question already in their organization, must send the reports.

The same is true for organization filters. Let's assume you belong to two organizations: OnlyHosts and OnlySwitches. If you set your organization filter to display only the elements in OnlySwitches and not in OnlyHosts, your certain reports display only the elements in OnlySwitches. If you send an Assets by Department report by e-mail, the information in the report displays information about the elements that are currently allowed through the organization filter. You can set the organization filter when you create an e-mail schedule. See ["Sending a Report by E-mail"](#) on page 343.

The following reports displays all information regardless of a user's organizations:

- Event Reports
- Application Reports
- Hosts Reports*
- Storage System Reports*
- Switch Reports*

*Reports that display information for only one element, such as Asset Summary, Details, Events, and Utilization reports.

Accessing Reporting

To access Reporting, click **Reports > Storage Essentials > Manage Reports**.

Viewing Reports

Collectors gather information for reports. If you stop a collector, its reports are not updated. You might want to verify a collector is running if you are having difficulty viewing a report or you believe the report is not displaying the latest information. See [“Viewing Report Collectors for an Element”](#) on page 340 for more information.

Reports in general contain no data or incomplete/incorrect data while reports are being materialized (materialized views are being populated).

IMPORTANT: The elements you see in the report are based on your organizations and which organizations are selected in the organization filter at the top of the page.

To view a report:

1. Access Reporting as described in [“Accessing Reporting”](#) on page 339.
2. In the middle pane, expand the tree and click the type of report you want.
The report is shown in the right pane.
3. To change the format of the report, select one of the following from the **Report Format** drop-down menu and then click the **Refresh View** button:
 - **HTML (Default)** - The software displays the report in a Web page by default.
 - **PDF** - The software displays the report in Adobe Acrobat, a good option if you need to print the report. The software assumes you already have Adobe Acrobat Reader installed on your computer.
To obtain a copy of Adobe Acrobat Reader, go to <http://www.adobe.com>.
 - **Excel** - The software displays the report in Microsoft Excel, providing you have a copy of Microsoft Excel already installed.
 - **XML** - The software display the report in the XML format.
4. To view a report in a new window, click the Open in new window option located next to the drop-down menu. Then, click the Refresh View button.
5. To send a report by e-mail, use the E-mail Schedule tab. See [“Sending a Report by E-mail”](#) on page 343 for more information.

Viewing Report Collectors for an Element

The management server uses collectors to gather information. The Collectors tab provides information about the collectors for a particular element.

To start collectors and view reports for an element:

1. Access the **Collectors** page by doing one of the following:
 - Clicking an element in Application Viewer, and then clicking the **Collectors** tab.
 - Double-clicking an element in System Manager, and then clicking the **Collectors** tab.
 - Clicking an element in Chargeback Manager, and then clicking the **Collectors** tab.

2. To change a collector's start time, modify the time and date entered in the **Next Scheduled Run** field. If you decide to change the start time, make sure the date is in the yyyy-mm-dd format with the time resembling a 24-hour clock. There should be a space between the date and the time, as shown below:

2003-08-20 09:41

After the collector runs, the value in this column is updated to the next time the collector will run.

3. To change how often the collector runs, type the number of minutes in the **Interval** field.

IMPORTANT: Do not make the interval too short. Running a collector too frequently uses up space on the management server and impacts its performance.

4. To enable the collector, click the **Start** button.
5. To stop a collector, click the **Stop** button.
6. To view a report, click its link. See "[Viewing Reports](#)" on page 339 for more information.

Refreshing a Report

If you want to view the latest information in a report, click the **Refresh Now** button on the Report Cache tab (**Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager). The management server gathers the latest information from the database and makes this information available to the reports.

The reports are refreshed every six hours by default. You can change the schedule for refreshing a report. See "[Refreshing the Report Cache](#)" on page 128.

If you are still seeing old information after you click the **Refresh Now** button, verify your database is being updated within the appropriate time frame for your organization. Some of the information contained in the database depend on collectors. Verify that the collectors are running at the appropriate intervals for your organization. Click the **Configuration** tab to access the configuration pages for the many types of collectors.

If you find you are still viewing old information regarding elements on the network, you may need to perform Discovery Data Collection. It is best to perform Discovery Data Collection at regular intervals. See "[Adding a Discovery Schedule](#)" on page 100.

Changing the Formatting of a Report

Reports are displayed in HTML format by default. You can change the formatting of the report by selecting one of the following options from the **Format** drop-down menu in Reporting and then clicking the **Apply** button:

- **HTML (Default)** - The software displays the report in a Web page by default.
- **PDF** - The software displays the report in Adobe Acrobat, a good option if you need to print the report. The software assumes you already have Adobe Acrobat Reader installed on your computer. To obtain a copy of Adobe Acrobat Reader, go to <http://www.adobe.com>.
- **Excel** - The software displays the report in Microsoft Excel, providing you have a copy of Microsoft Excel already installed.
- **XML** - The software display the report in the XML format.

Opening a Report in a New Window

Use this feature to view two or more reports simultaneously.

To view a report in a new window:

1. Select the **Open in new Window** option.
2. Click the **Apply** button.
A new window opens, and it displays just the report.
3. You can arrange the windows on the screen so you can view one report in one corner and another report in a different corner. You can also toggle between the two reports, depending on your operating system.

Maximizing the Screen Space for a Report

In some instances you might need to maximize the screen space to view a report. This can be done by hiding the middle and left pane, as described below:

1. To hide the middle pane, click the border between the tree listing the reports and the main pane and drag it to the left side of the page, as shown in the following figure.




Figure 59 Hiding the Middle Pane

2. To display the middle pane, drag the border to the right.

Filtering Data in Global Reports

IMPORTANT: Depending on your license, Global Reporter may not be available. See the “List of Features” to determine if you have access to global reports. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials). If your license lets you access global reports but you cannot access them, contact your system administrator to make sure your role lets you view and/or set up global reports.

You can filter the data in global reports so you only see the data gathered from certain sites and assigned to certain organizations. For example, assume your Global Reporter server gathers data from 10 sites throughout the world. You can filter the Global Reporter reports so you see data gathered from just one site.

If you do not see information in your reports, verify you have global reporting set up correctly. See “[Setting Up Global Reporter](#)” on page 130.

To filter data global reports:

1. Access Reporting as described in "[Accessing Reporting](#)" on page 339.
2. Expand the **Global** node.
3. Select a global report.
4. In the right pane, click the **Filter Data** button.
5. Expand the **Sites and Organizations** node.
6. In the Filtering window, select the sites you want to view.
7. Expand the node of each selected site to view its organizations.
8. Select the organizations containing the elements you want to view. To view all organizations at a site, select the **Everything** option.
9. Click **OK**.

The filter applies to all global reports.

Sending a Report by E-mail

You can send by e-mail an attached report in PDF, XML or Microsoft Excel format. If you want to send reports by e-mail on a regular basis, set up an e-mail schedule for the report, as described in "[Adding an E-mail Schedule for a Report](#)" on page 344.

Keep in mind the following:

- Before you can e-mail a report, you must set up e-mail notification, as described in the topic, "[Setting Up E-mail Notification](#)" on page 99.
- The elements in the report you send are based on your organizations and which organizations are selected in the organization filter at the top of the page.
- Send your reports soon after a report cache refresh. The reports display data that is in the report cache. If the report cache contains old data, the reports you send by e-mail will also show old data. The reports are refreshed every six hours by default. See "[Scheduling a Report Cache Refresh](#)" on page 129.

To send a report by e-mail:

1. Access Reporting as described in "[Accessing Reporting](#)" on page 339.
2. Expand the tree in the middle pane, and click the report you want to send by e-mail.
3. Click the **E-mail Report** button.
You are told the E-mail server is not enabled if you have not set up e-mail notification. You must set up e-mail notification before you can send e-mail a report. See "[Setting Up E-mail Notification](#)" on page 99.
4. In the top field, type the recipient's e-mail address.
The software verifies the address entered has a correct form. To send multiple addresses, separate each address with a comma (,), for example:
`john.example@appiq.com,jerry.example@appiq.com`
5. From the **Format** drop-down menu, select one of the following formats:
 - **PDF** - Requires the use of Adobe Acrobat, which can be downloaded for free from <http://www.adobe.com>.
 - **Excel** - Requires the use of Microsoft Excel.

- **XML** - Requires the user has an understanding of XML.
6. (Optional) Modify the subject and message.
 7. Click the **OK** button.
- The report is sent.

Managing E-mail Schedules for Reports


This section describes the following:

- ["Adding an E-mail Schedule for a Report"](#) on page 344
- ["Editing an E-mail Schedule for a Report"](#) on page 346
- ["Deleting E-mail Schedules for a Report"](#) on page 348
- ["Viewing E-mail Schedules for a Report"](#) on page 349

Adding an E-mail Schedule for a Report

You can add an e-mail schedule so that a user receives an attached report on a regular basis. The report can be in the form a PDF, XML or Microsoft Excel document.

Keep in mind the following:

- Before you can add an e-mail schedule, you must set up e-mail notification, as described in ["Setting Up E-mail Notification"](#) on page 99.
- Schedule to send your reports soon after a report cache refresh. The reports display data that is in the report cache. If the report cache contains old data, the reports you send by e-mail will also show old data. The reports are refreshed every six hours by default. For example, assume you added an e-mail schedule that sends a report daily at 7 a.m. Also, assume you scheduled your report cache refreshes to take place daily at 8 a.m. Your reports will most likely show outdated data. It would make more sense to schedule your report cache refresh at 7 a.m. and then schedule to send your reports soon afterwards. See ["Scheduling a Report Cache Refresh"](#) on page 129.
- The management server service must be running for users to receive e-mail notification.
- Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager. Then, click the **Scheduled Deliveries** tab at the top of the screen.
- The elements in the report you send are based the organizations selected from the Organization Filters tab. For example assume you belong to two organizations: OnlyHosts and OnlySwitches. If you select OnlySwitches and not OnlyHosts in the Organization Filters tab (accessible by clicking **Add E-mail Schedule**), the user only receives information about the elements in the OnlySwitches organization. This is still true even if you change your organization filtering at the top of the page (.

To add an e-mail schedule:

1. Access Reporting as described in ["Accessing Reporting"](#) on page 339.
2. Expand the tree in the middle pane, and click the report you want to send at a scheduled time.

3. When the report is displayed in the right pane, click the **Scheduled Deliveries** tab in the right pane.
4. Click the **Add E-mail Schedule** button.
5. Verify the **Properties** tab is displayed.
6. In the **to** field, type the recipient's e-mail address.
The software verifies the address entered has a correct form. To send multiple addresses, separate each address with a comma (,), for example:
john.example@appiq.com,jerry.example@appiq.com
7. In the **Subject** field, type a subject for the e-mail messages you plan to send.

NOTE: Provide the name of the report in the subject field so users can distinguish the message from others.

8. In the **Message** field, type a message describing the report.
If you are e-mailing reports in bulk, you might want to let users know the e-mail is being sent by an automated process. You might also want to provide an e-mail address for users to provide feedback, for example:

This e-mail and its attached report are generated automatically. If you would like to change how often the report is sent to you or you want to be taken off the list, please contact username@companyname.com.
9. From the **Format** drop-down menu, select one of the following formats:
 - **PDF** - Requires the use of Adobe Acrobat, which can be downloaded for free from <http://www.adobe.com>.
 - **Excel** - Requires the use of Microsoft Excel.
 - **XML** - Requires the user has an understanding of XML.
10. In the **Time to Run** field, type the time you want to send the report. This time must be entered in the 24-hour format. For example, if you want a report sent at 2:15 p.m., you would type 14:15 in the **Time to Run** field.
11. Select one of the following options to determine how frequently you want to send the report.
 - **Daily** - If you selected daily, select how frequently you want the management server to send the report.
 - **Everyday** - The report is sent everyday.
 - **Weekday** - The report is sent only Monday through Friday.
 - **Every x days** - Fill in the frequency you want the report to be sent. For example, if you enter 15, the report is sent every 15 days.
 - **Weekly** - If you selected weekly, use the **Frequency** drop-down menu to select the day of the week on which you want the report sent.
 - **Monthly** - If you selected monthly, select the time during the month you want the report sent.
 - To send the report on the first or last day of the month, select the first option. Then, select **First** or **Last** from the drop-down menu.

- To send the report on a specified day during the month, select the second option. Then, type the day on which you want the report sent. If you type a day that is not in the month, for example 30 for February, the report is sent on the last day of the month.

12. Click **Next.**

The Organization Filters tab is not available for all reports. If an Organization Filters tab is not available, you see a **Finish** button instead of a **Next** button. Click the **Finish** button and skip the rest of the steps in this procedure.

13. Select the organizations containing the elements you want used in the report. If you find all organizations selected, deselect the organizations containing the elements you do not want displayed in the report.

Organizations you belong to and their children are displayed. Only elements belonging to selected organizations are displayed in the report you send. For example, assume you selected OnlyHosts, as shown in the following figure. The user receiving the report would only see data about elements in OnlyHosts. Information about elements in other organizations would not be displayed.



Figure 60 Selecting Organizations Used in This Report


14. Click **Finish.**

The schedule is created.

Editing an E-mail Schedule for a Report

IMPORTANT: Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager. Then, click the **Scheduled Deliveries** tab.

To edit an e-mail schedule for a report:

1. Access Reporting as described in "Accessing Reporting" on page 339.
2. Expand the tree in the middle pane, and click the report you want to send at a scheduled time.
3. When the report is displayed in the right pane, click the **Scheduled Deliveries** tab in the right pane.
4. Under the Edit column, click the **Edit** () button.
5. Verify the **Properties** tab is displayed.
6. In the top field, type the recipient's e-mail address.

The software verifies the address entered has a correct form. To send multiple addresses, separate each address with a comma (,), for example:

john.example@appiq.com, jerry.example@appiq.com

7. In the **Subject** field, change the subject of the e-mail.
8. In the **Message** field, change a message describing the report.

If you are e-mailing reports in bulk, you might want to let users know the e-mail is being sent by an automated process. You might also want to provide an e-mail address for users to provide feedback, for example:

This e-mail and its attached report are generated automatically. If you would like to change how often the report is sent to you or you want to be taken off the list, please contact username@companyname.com.

9. From the **Format** drop-down menu, select one of the following formats:
 - **PDF** - Requires the use of Adobe Acrobat, which can be downloaded for free from <http://www.adobe.com>.
 - **Excel** - Requires the use of Microsoft Excel.
 - **XML** - Requires the user has an understanding of XML.
10. In the **Time to Run** field, type the time you want to send the report. This time must be entered in the 24-hour format. For example, if you want a report sent at 2:15 p.m., you would type 14:15 in **Time to Run** field.
11. Select one of the following options to determine how frequently you want to send the report.
 - **Daily** - If you selected daily, select how frequently you want the management server to send the report.
 - **Everyday** - The report is sent everyday.
 - **Weekday** - The report is sent only Monday through Friday.
 - **Every x days** - Fill in the frequency you want the report to be sent. For example, if you enter 15, the report is sent every 15 days.

- **Weekly** - If you selected weekly, use the **Frequency** drop-down menu to select the day of the week on which you want the report sent.
- **Monthly** - If you selected monthly, select the time during the month you want the report sent.
 - To send the report on the first or last day of the month, select the first option. Then, select **First** or **Last** from the drop-down menu.
 - To send the report on a specified day during the month, select the second option. Then, type the day on which you want the report sent. If you type a day that is not in the month, for example 30 for February, the report is sent on the last day of the month.

12. Click **Next**.

The Organization Filters tab is not available for all reports. If an Organization Filters tab is not available, you see a **Finish** button instead of a **Next** button. Click the **Finish** button and skip the rest of the steps in this procedure.

13. Select the organizations containing the elements you want used in the report. If you find all organizations selected, deselect the organizations containing the elements you do not want displayed in the report.

Organizations you belong to and their children are displayed. Only elements belonging to selected organizations are displayed in the report you send. For example, assume you selected OnlyHosts, as shown in the following figure. The user receiving the report would only see data about elements in OnlyHosts. Information about elements in other organizations would not be displayed.




Figure 61 Selecting Organizations Used in This Report

14. Click **Finish**.

Deleting E-mail Schedules for a Report

IMPORTANT: **Important:** Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager. Then, click the **Scheduled Deliveries** tab.

To delete an e-mail schedule:

1. Access Reporting as described in "[Accessing Reporting](#)" on page 339.
2. Expand the tree in the middle pane, and click the report corresponding to the e-mail schedule you want to delete.
3. When the report is displayed in the right pane, click the **Scheduled Deliveries** tab in the right pane.
4. Click the  button corresponding to the e-mail schedule you want to remove.

Viewing E-mail Schedules for a Report

IMPORTANT: Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager. Then, click the **Scheduled Deliveries** tab.



To view the E-mail schedules assigned to a report.

1. Access Reporting as described in "[Accessing Reporting](#)" on page 339.
2. Expand the tree in the middle pane, and click the report corresponding to the e-mail schedules you want to view.
3. When the report is displayed in the right pane, click the **Scheduled Deliveries** tab in the right pane.
Information about the e-mail schedules for that report are displayed.

Table 57 Viewing E-mail Schedules for a Report

Column Name	Description
Recipient	The person who receives the report.
Subject	The subject of the e-mail, brief summary of what it is about.
Format	The format of the report sent: <ul style="list-style-type: none">• PDF• Microsoft EXCEL• XML

Table 57 Viewing E-mail Schedules for a Report (continued)

Column Name	Description
Edit	Click the  button to edit a schedule of the report. See "Adding an E-mail Schedule for a Report" on page 344 and "Editing an E-mail Schedule for a Report" on page 346 for information about the options displayed in this window.
Delete	Click the  button to remove the corresponding schedule.

Creating Custom Reports

This section describes the following:

- ["About Creating Custom Reports"](#) on page 350
- ["Configuring Report Designer to Work with the Management Server"](#) on page 351
- ["Designing Custom Reports"](#) on page 352
- ["Integrating Custom Reports"](#) on page 361
- ["Detailed Schema Information"](#) on page 363
- ["Views from Previous Releases"](#) on page 396

IMPORTANT: You must install Report Designer before you can create custom reports. Obtain a copy of Report Designer from your sales professional. Follow the installation instructions that accompany it.

About Creating Custom Reports

To create customized reports, you need a program for creating reports, such as Report Designer. Use Report Designer to create customized reports on the management server. Report Designer links to the database of the management server, so you can view real-time data in your customized reports. Once you are satisfied with the customized reports, you can merge them onto the management server so they are accessible from the management server console.

The following figure shows how Report Designer fits into the reporting architecture. Use Report Designer to create the reports. Then, deploy the reports on the management server.

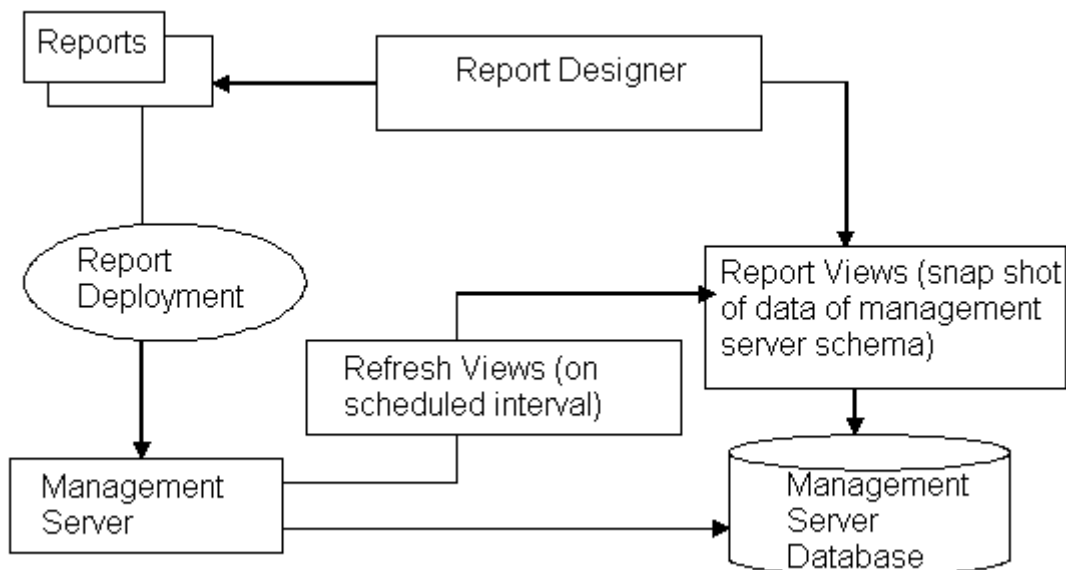


Figure 62 Report Architecture

When you create reports, you can use pre-existing schema materialized views. A materialized view is a snapshot of data, from the database, created from a query. Report Designer refers to these materialized views as tables. Materialized views are refreshed based on its collector's schedule for obtaining the latest data. The default refresh time is every six hours. The current and deprecated materialized views are provided at the end of this chapter.

Configuring Report Designer to Work with the Management Server

IMPORTANT: The steps in this section are for configuring Report Designer version 6.0.

You can use data from the management server to create your custom reports. These steps assume you have installed Report Designer on a Microsoft Windows computer that can access the `C:\oracle\ora92\jdbc\lib` directory on the management server.

To use the management server with Report Designer:

1. Add the following class path to Report Designer. Refer to the documentation accompanying Report Designer for more information.

`C:\oracle\ora92\jdbc\lib\classes12.jar`

where C:\oracle\ora92 is the directory containing classes12.jar.

2. Select File > New Catalog.

A catalog is a repository for reports. The catalog and the reports that are based on it must be in the same directory for the report to run. This is because the catalog contains the object definitions that are used by the reports in the catalog.

3. In the name field, type a name for the catalog, such as custom.cat.

4. Save the catalog in a directory especially for it. This way you can work on the reports remotely and then easily move them when you are ready to integrate them with the product.

5. Click OK.

When you are done with creating the catalog, the Catalog Browser window appears. If you do not see the Catalog Browser window, open the catalog by selecting **File > Open Catalog**. Select the catalog you want to open, and then click the **Open** button.

6. In the Catalog Browser window, expand the Default node in the tree. Then, right-click Connection and select New Connection.

7. Verify that the OracleOraHome92TNSListener service is running on the management server. If you want to view live data in your custom reports, the management server does not need to be running; however, the Oracle database for the management server does need to be running.

8. In the Get JDBC Connection Information window do the following:

a. Deselect the Use ODBC Data Source option.

b. Select the JDBC Driver option.

c. Type the following for the JDBC driver: `oracle.jdbc.driver.OracleDriver`

d. Type the following in the JDBC URL field:

```
jdbc:oracle:thin:@HostIP/DNS:1521:APPIQ
```

where HostIP/DNS is the host IP address or DNS name of the host running the management server

If Report Designer is running on the same computer as the management server, you can use localhost for the DNS name, as shown in the following example:

```
jdbc:oracle:thin:@localhost:1521:APPIQ
```

e. Type the following in the User Name field: `Report_User`

This is the user name that is used to access the schema view in the management server database. This user has read privileges only for the schema views.

f. Type the following in the Password field: `appiq`

9. Click OK.

Report Designer searches for JDBC driver.

You may have entered incorrect path information if Report Designer cannot find the JDBC drivers.

10. Select APPIQ_SYSTEM under the schemas section.

The Tables pane becomes populated. These are all the tables you can use to create the reports. It is best to select as many tables as possible rather than too few. If you are not too sure as to which tables you may be using, you may want to select them all.

11. Select the tables and then click **Add**. Click the **Done** button when you have finished adding tables.

The tables populate the Catalog Browser window.

12. Click **File > Save Catalog** to save the catalog.

Designing Custom Reports

IMPORTANT: This section assumes you have already installed and configured Report Designer and integrated it with the management server.

These steps assume you are running Report Designer 6.0, and you have the online help for Report Designer installed. The instructions provide general information about using Report Designer. For additional information, refer to the online help for Report Designer. The management server only supports Report Designer running on Microsoft Windows.

NOTE: You cannot create custom reports for NetApp filers.

Creating Standard Reports

The following steps are for Microsoft Windows and for Report Designer 6.0:

1. Open Report Designer.
2. If the Choose Report screen does not appear, select **File > New** in Report Designer.

3. Click the **Standard Report** icon and then click **Create**.

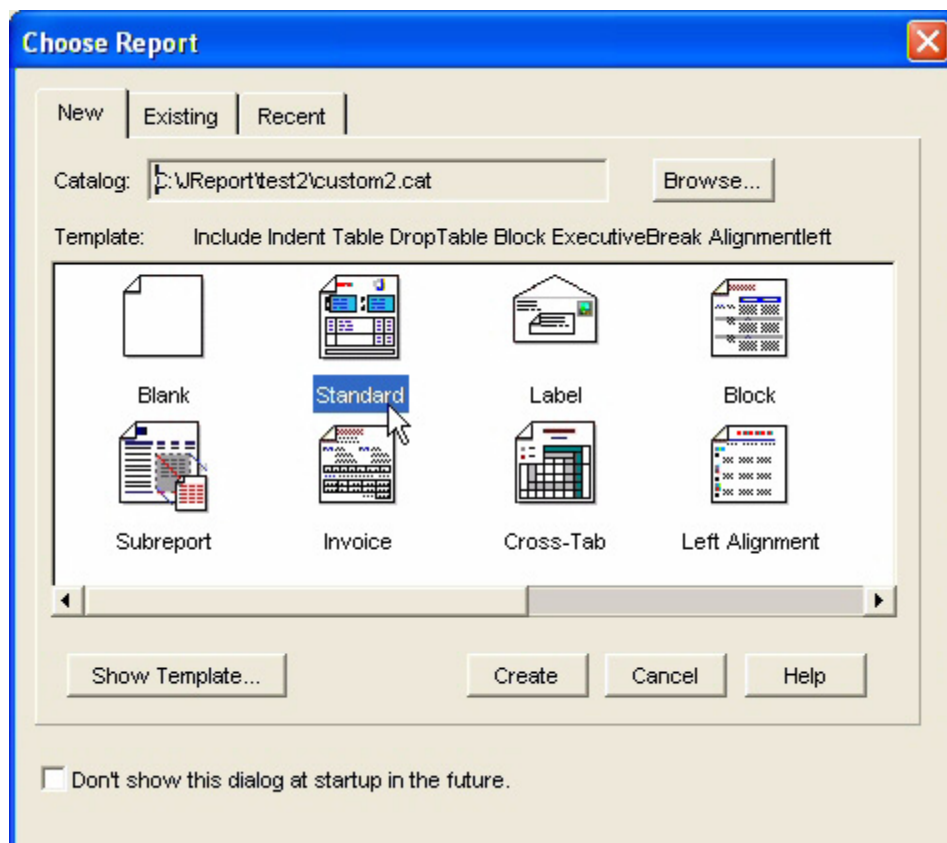


Figure 63 Choosing a Standard Report

4. Select the data you want in your report by selecting the corresponding materialized views (tables) displayed in the Data tab. Use the Report Categories table in ["About Creating Custom Reports"](#) on page 350 as a guideline. For example, in the following figure, the MV_APPLICATIONVW table has been selected. System application data will be made available to the report, according to the Report Categories table. You can, however, specify

you do not want all data displayed in the report. When you are done, click **Next**. To find a definition of the listings in a table, see "[Detailed Schema Information](#)" on page 363.

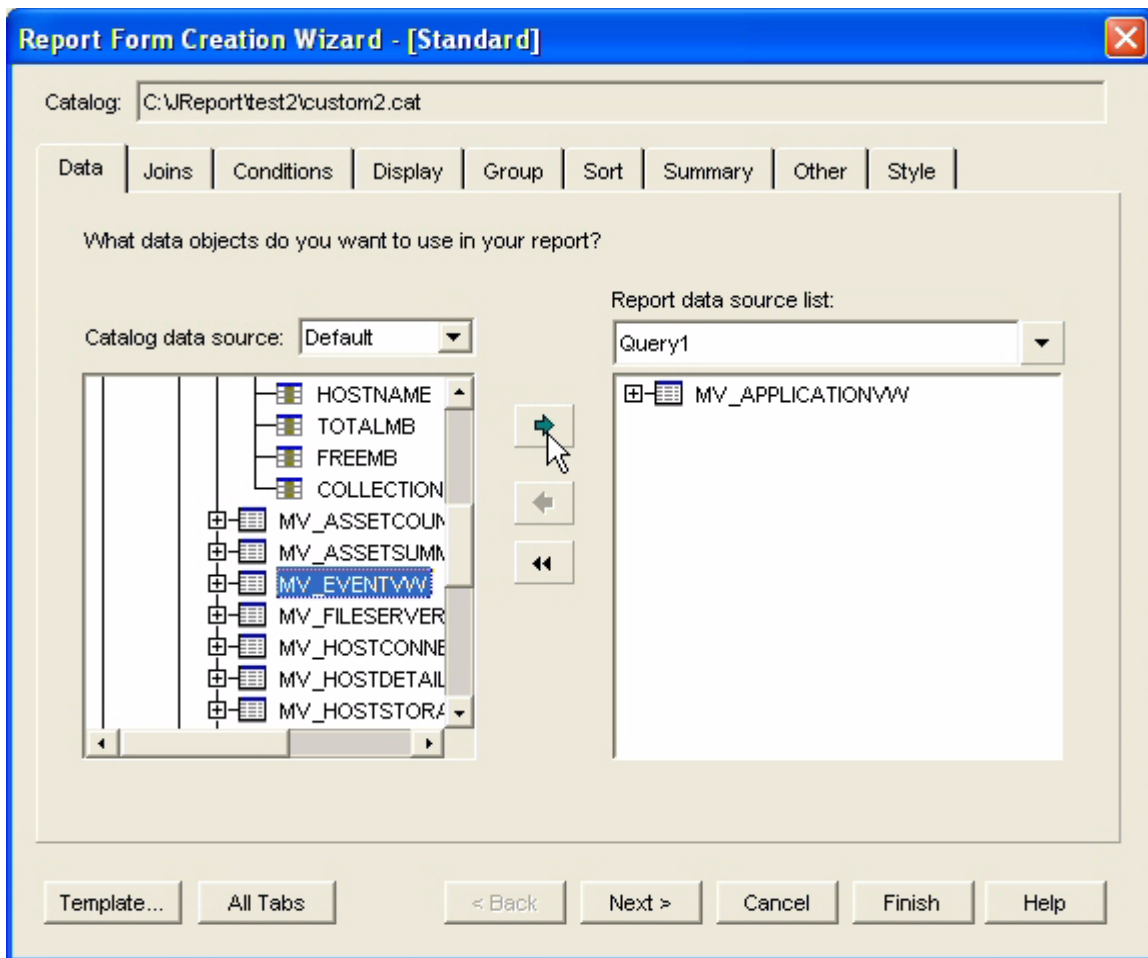


Figure 64 Adding Tables for a Standard Report

5. When you are asked if Report Designer will create a new query, click **OK**.
6. If you selected more than one materialized views (tables) in the Data tab, you need to link common search criteria, such as the `DEVICE_ID` in one table to the `DEVICE_ID` of another table. Sometimes the search criteria will have different terminology. For example, in the following figure, `MV_EVENTVW_DOMAINNAME` is linked with `DOMAIN NAME`. If a search criteria is cut off in a table, you can expand the size of the table by clicking a table border and dragging it to the appropriate position.

To find a definition of a search criteria, see "[Detailed Schema Information](#)" on page 363.

Refer to the online help for Report Designer for more information. When you are done, click **Next**.

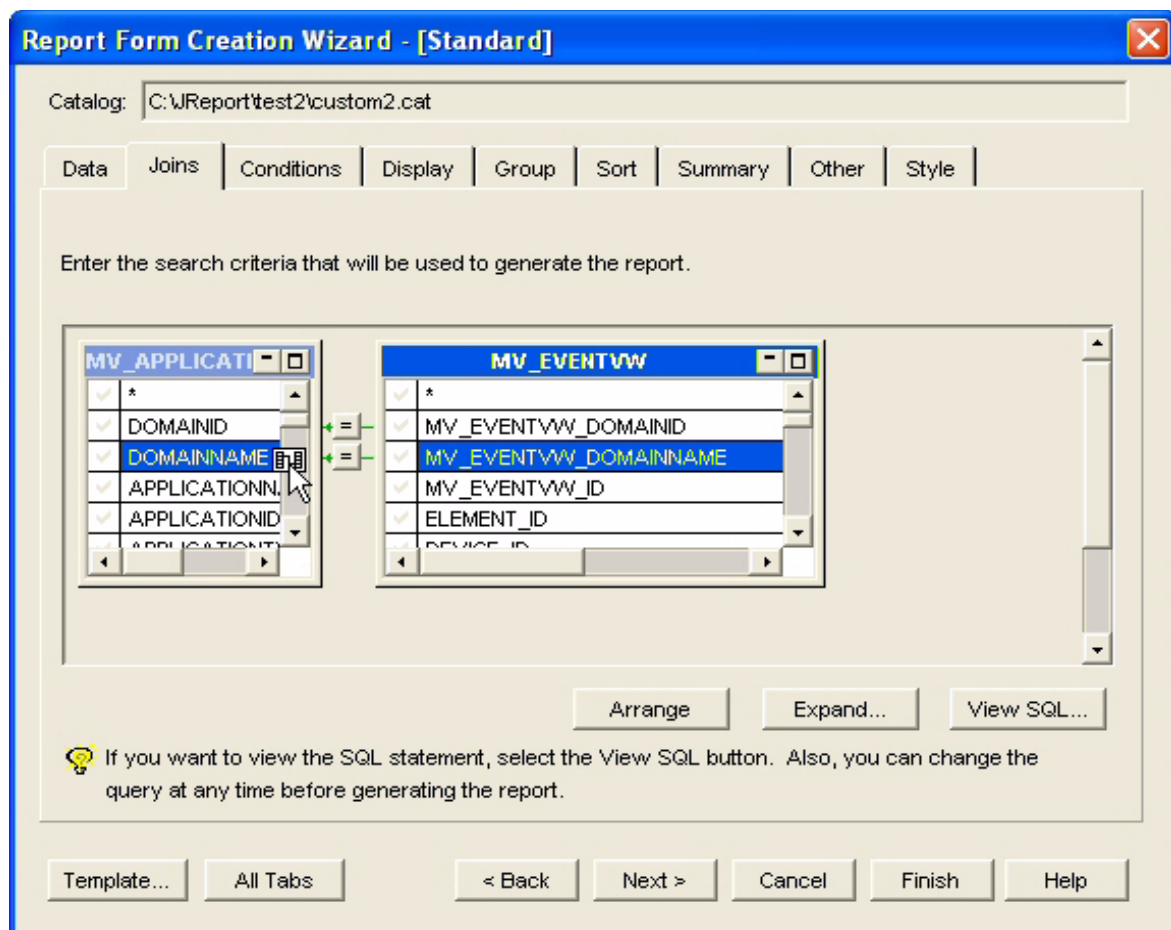


Figure 65 Linking Common Data in Tables for a Standard Report

7. Enter search criteria that will be used to generate the report. For example, let's assume you want the report to display information only about Oracle applications. You would enter a

search criteria that tells Report Designer to display data from Oracle applications. When you are done, click **Next**.

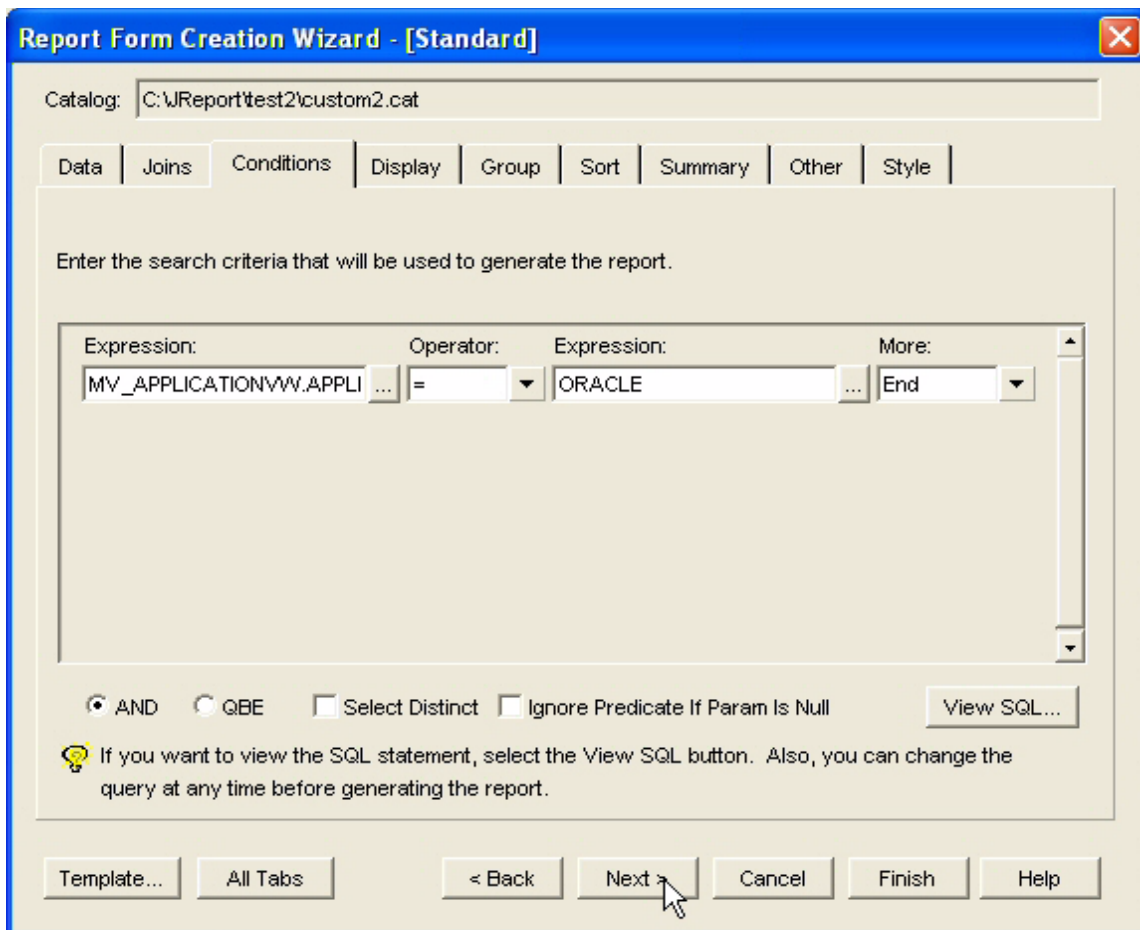


Figure 66 Creating Search Criteria for Standard Reports

8. Select the data you want displayed in the report. Click the data source in the left and then click the arrow pointing right. Its report field appears in the Report fields pane. When you are done, click **Next**.

The order of the report fields in the Report fields pane determines their sequence in the report. For example, a report field at the top of the list in the Report fields pane will appear in the far left column in the report. Likewise, a report field at the bottom of the list in the Report fields pane will appear in the far right column in the report. The column heading for the data is determined

by the text in the AutoLabel column in the Report fields pane. To find a definition of the listings in a table, see “[Detailed Schema Information](#)” on page 363.

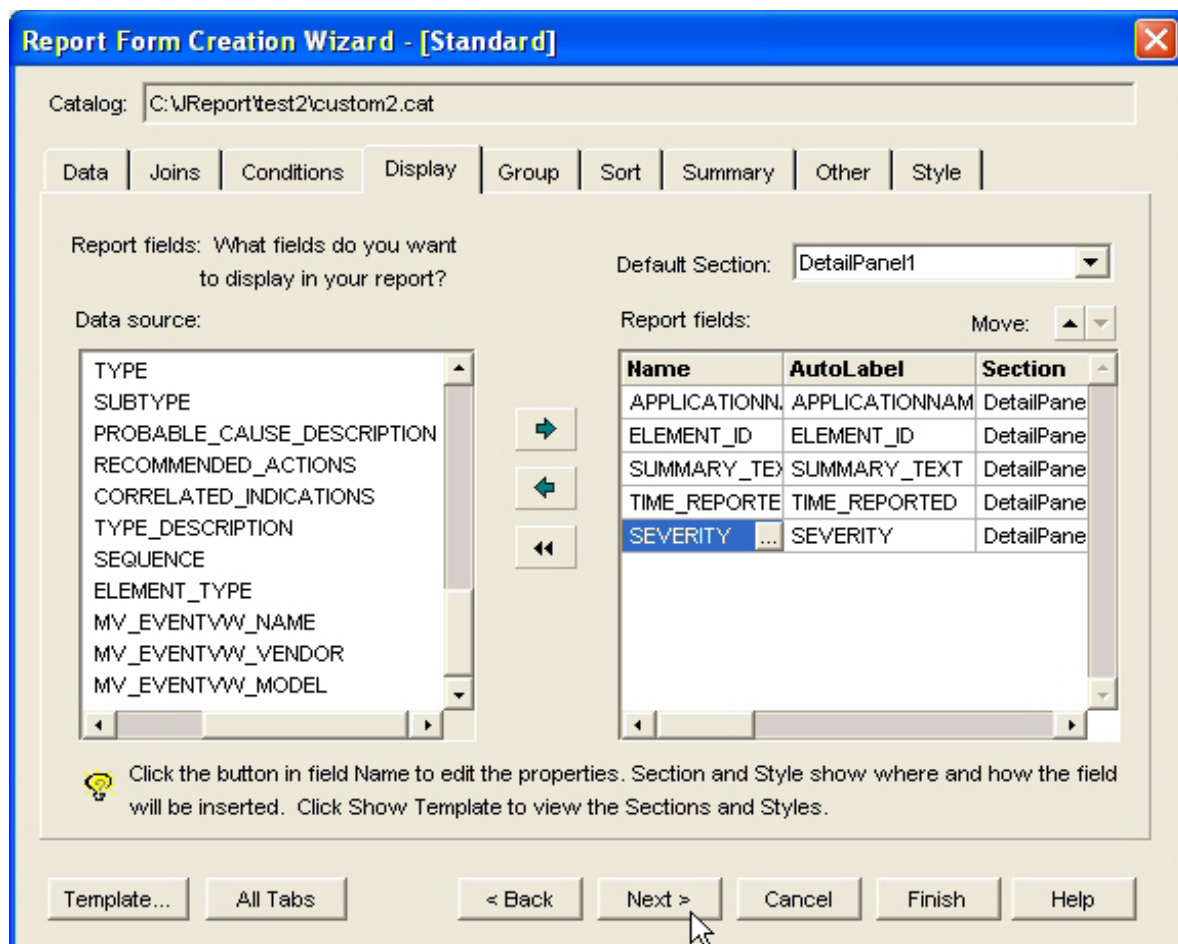


Figure 67 Deciding Which Data Should Appear in the Report

9. Select the fields in the left pane in the order you want them sorted in your report. Then, click the arrow pointing right. When you are done, click **Next**.

For example, in the following figure, information in the report will first be sorted by an application name. Let's assume you then selected APPLICATIONID. Applications would first be sorted by application name and then by their application ID.

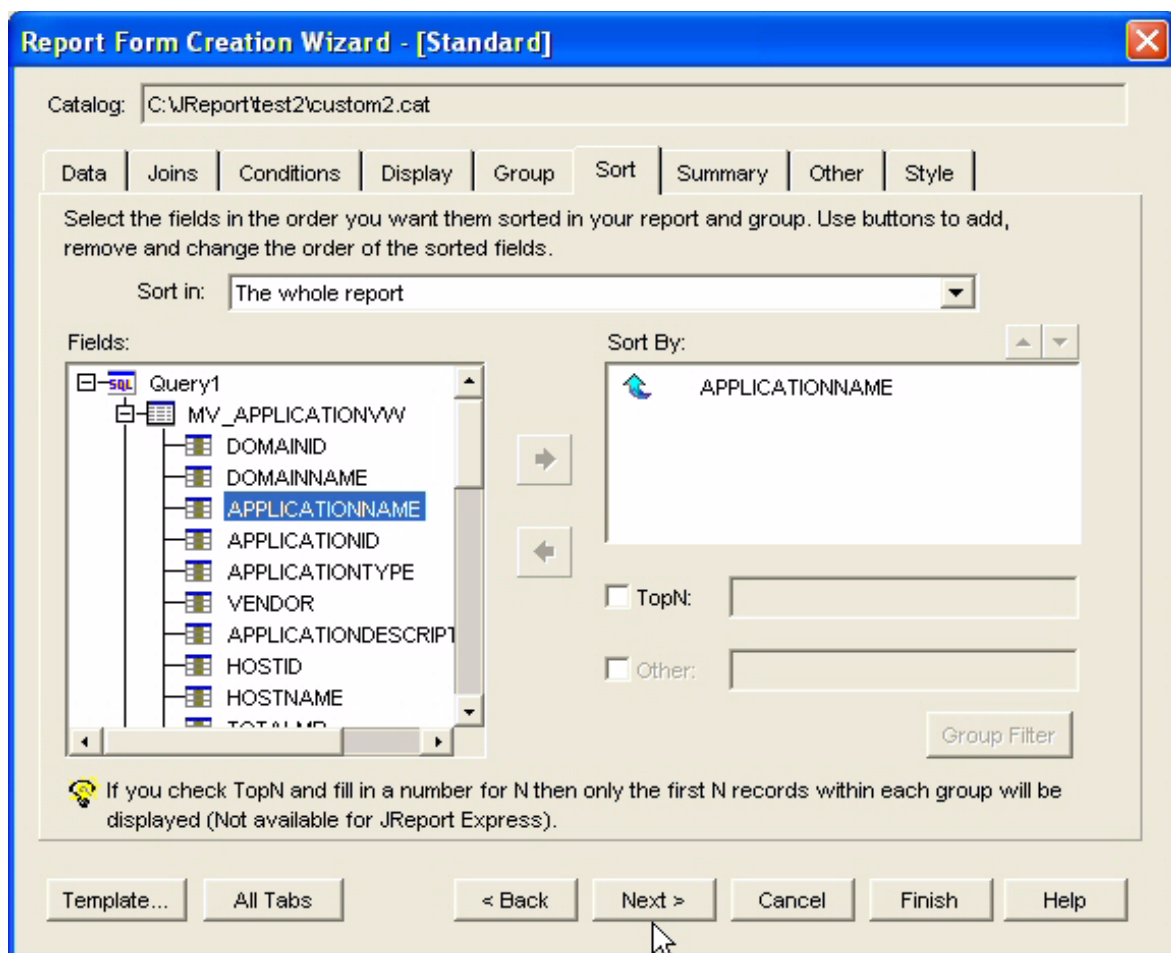


Figure 68 Sorting Information in the Report

10. Use the Style tab to determine the layout of the report. When you are done, click **Finish**.

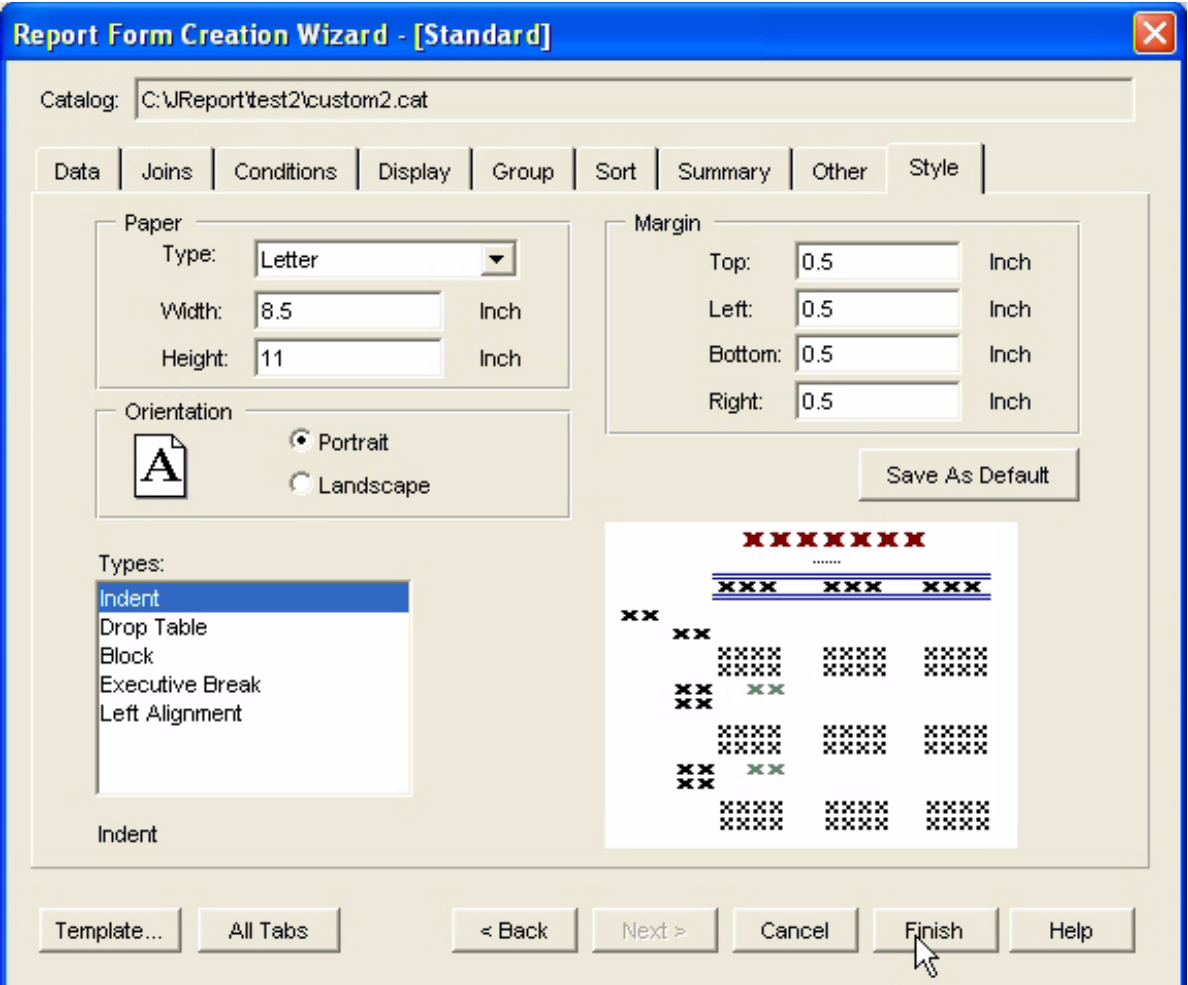


Figure 69 Selecting the Layout of the Report

The report template is displayed. You will not see any data reported, only placeholders, as shown in the following figure.

Report Title				
APPLICATIONNAME	ELEMENT_ID	SUMMARY_TEXT	TIME_REPORTED	SEVERITY
XXXXXXXXXX				
XXXXXXXXXX	###.DD	XXXXXXXXXX	MM/dd/yy	###.DD

Figure 70 Report Template Displayed

11. To view the report with its data, click the **View** tab.

12. The database for the management server must be running on the management server to be able to view active data in the report. Verify that the OracleOraHome92TNSListener service is running on the management server.

If you do not see any text, verify that the management server has collected this data. See which tables in your custom report map to pre-existing reports. Use the table in “[About Creating Custom Reports](#)” on page 350 as a guideline. Then, access Reporting on the management server and verify the corresponding reports are displaying information.

If you are running Report Designer remotely from the management server and you cannot view active data in your reports even with the Oracle database running, verify you did the following when you installed Report Designer.

- Save classes12.jar from the management server on your computer and pointed this file in the class path when you installed Report Designer. If so, the file is listed in [Report Designer installation directory]\jreport.bat.
- Verify the connection information for the catalog is correct. Open the catalog and expand the Connection node. Verify the IP address/DNS name listed is correct.

If you are still having problems, verify the network from your computer to the management server is stable. Try accessing the management server console from your computer as a test.

13. If you want to view live data in your custom reports, the management server does not need to be running; however, the Oracle database for the management server does need to be running.

14. Click **File > Save Template** to save the report.

15. Refer to the online help for Report Designer for information on how to design the reports.

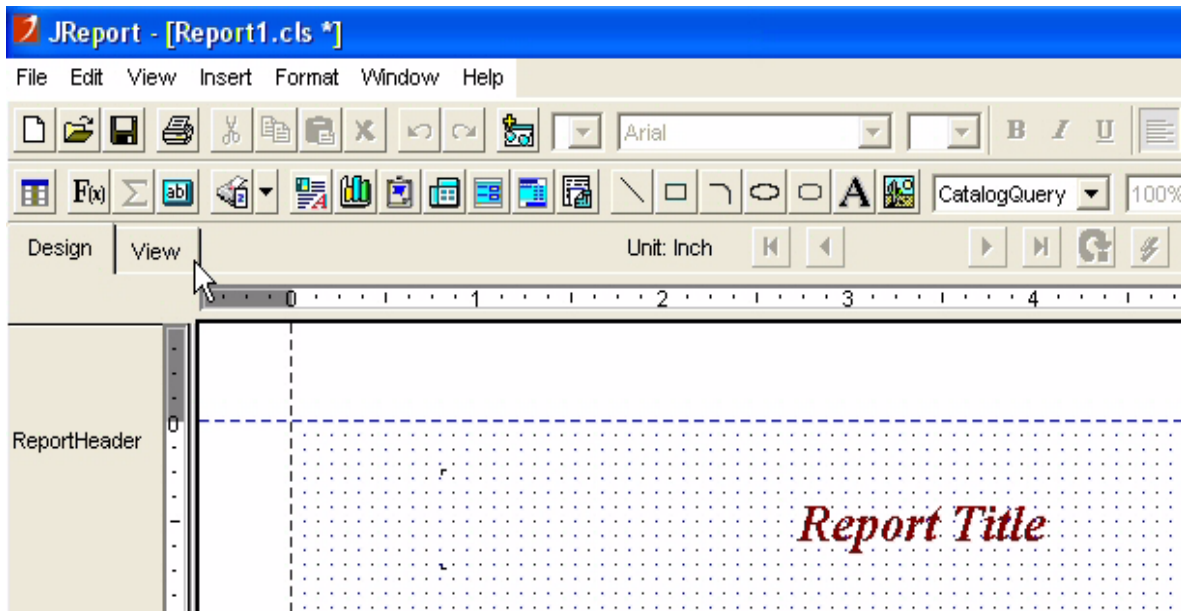


Figure 71 Clicking the View Tab

Integrating Custom Reports

After you are satisfied with your custom report, you must integrate it with the management server so that other users can access the report. You must first deploy the custom report to the management server. Then, you must integrate the report so that it is accessible from Reporting. Custom reports appear with their own icon in the reports tree.

IMPORTANT: You can have only one catalog (*.cat) file to a directory. If you have more than one catalog, create a subdirectory under
%JBOSS4_DIST%\server\appiq\reports\custom. For example, assume you have two reports that use a different catalog, you would create two subdirectories under reports\custom. The support files and catalog for one report would go into one subdirectory (reports\custom\subdirectory1), and the support files and catalog for the other report would go into another subdirectory (reports\custom\subdirectory2).

To integrate custom reports:

1. Deploy the custom report files (CLS, CAT and all other files) from Report Designer. Use Report Designer's deploy catalog mechanism to transfer the files. Create a directory called `custom` under %JBOSS4_DIST%\server\appiq\reports. Then, deploy the catalog to the following directory:

%JBOSS4_DIST%\server\appiq\reports\custom

Refer to the documentation accompanying Report Designer for more information.

2. Create a node for your custom reports. This node appears in the tree in Reporting.
 - a. Create the directory `%JBOSS4_DIST%\server\appiq\reports\customTreeNodes` if it doesn't already exist
 - b. Create an XML file in that directory for your tree node, for example, `custom.xml`. The management server uses this file to determine where to put the node for custom reports in the tree.
 - c. Use a text editor, such as Notepad, to open the XML file you created in the previous step. Enter the following into the XML file:

```
<NODE NAME="CUSTOM" LABEL="Custom" PARENT="SYSTEM"/>
```

where

- `NAME` is how reports and other tree nodes refer to this node.
- `LABEL` is the label that appears in the user interface for the node.
- `PARENT` is the name of the parent node in the tree (optional). In this case, the Custom report appears under the SYSTEM node in the tree in Reporting.

3. To populate the tree in Reporting with your custom reports:
 - a. Create the directory `%JBOSS4_DIST%\server\appiq\reports\definitions\custom` if it doesn't already exist
 - b. Create an XML file in that directory for your report, for example: `CustomHostReport.xml`
 - c. Open the file you created in the previous step in a text editor, such as Notepad. Add the following information to the file, modifying it for your settings. Refer to files in `%JBOSS4_DIST%\server\appiq\reports\definitions` if you need additional samples.
 - d. Make sure you use an ID that is not used by any of the existing reports. You have to specify a title to appear in the tree and the file names for the CLS and CAT files, as shown in the following example.


```
<REPORT ID = "9098"
  TITLE = "Custom Host Connectivity"
  FILE_NAME = "//definitions/custom/Host_Connectivity.cls"
  CATALOG_FILE = "//definitions/custom/host.cat"
  SUPPORTS_ORGANIZATION_FILTERS = "true">
  <TREE_NODE NAME = "CUSTOM" />
  <ELEMENT_FUNCTION_CALLBACK NAME = "isHost" />
</REPORT>
```

where

- `REPORT ID` is the unique ID for the report.
- `TITLE` is the name you want to appear for the report in the tree in Reporting.
- `FILE_NAME` is the file name of the CLS file for the report and its path. In the example, the CLS file is in the `%JBOSS4_DIST%\server\appiq\reports\definitions\custom` directory.

- `CATALOG_FILE` is the file name of the CAT file for the report and its path. In the example, the CAT file is in the `%JBOSS4_DIST%\server\appiq\reports\definitions\custom` directory.
- `SUPPORTS_ORGANIZATION_FILTERS` if this property is set to true, the report supports organization filtering. If this property is set to false, the Organization Filters tab does not appear. If you leave omit `SUPPORTS_ORGANIZATION_FILTERS` the management server assumes you want to the report to support organization filtering.
- `TREE_NODE_NAME` is the name of the tree node you want the report to appear under. Use the node name of the file you created in `%JBOSS4_DIST%\server\appiq\reports\customTreeNodes`, not its file name.
- `ELEMENT_FUNCTION_CALLBACK_NAME` refers to the type of information you want to obtain from the report. For example if you want to obtain information from all hosts, enter `ishosts`. See the following table for a general listing of the element function callback names you should use. If you do not see your element listed, refer to the existing report definitions for the element located in `%JBOSS4_DIST%\server\appiq\reports\definitions`.

4. Restart the AppStorManager service.

The custom reports are displayed in the reports tree with a  icon.

Detailed Schema Information

This section provides information about the current materialized views. If you created reports in previous releases, you are most likely using old views. The old views are still supported; however, it is strongly suggested you use the new views in this section when creating reports. See “[Views from Previous Releases](#)” on page 396 for a mapping between the old and new views.

Table 58 Description of the Report Views

Materialized View (Tables)	Description
MVC_ORGANIZATIONVW	Provides information about organizations. See Table 74 on page 377.
MVC_ORGRELATIONVW	Provides information about the relationships in an organization, such as the parent and child identifiers in an organization. See Table 75 on page 377.
MVC_OPTIONALTABLEVW	Provides information about the names and values of the optional value. See Table 80 on page 379.
MVC_ASSETSUMMARYVW	Provides summary information about assets.
MVCA_BU_MASTERSERVERSUMMARY	Provides summary information about the master server. See Table 102 on page 391.

Table 58 Description of the Report Views (continued)

Materialized View (Tables)	Description
MVCA_BU_MEDIASERVERSUMMARY	Provides summary information about media servers. See Table 103 on page 392.
MVCA_BU_CLIENTSUMMARY	Provides summary information about clients. See Table 104 on page 392.
MVCA_BU_MEDIASUMMARY	Provides summary information about the media. See Table 105 on page 393.
MVCA_BU_JOBSUMMARY	Provides summary information about jobs. See Table 106 on page 394.
MVCA_BU_LIBRARYSUMMARY	Provides summary information about libraries. See Table 107 on page 395.
MVC_HOSTSUMMARYVW	Provides summary information about hosts. See Table 59 on page 366.
MVC_APPLICATIONSUMMARYVW	Provides summary information about applications. See Table 88 on page 385.
MVC_UNITACCESSVW	Provides information about unit access, such as access mode, host group, host group name, and host group modes. See Table 89 on page 386.
MVCA_DBAPPCAPACITYVW	Provides capacity information for a supported database application. See Table 90 on page 386.
MVCA_EXCHANGEAPPCAPACITYVW	Provides capacity information for Microsoft Exchange. See Table 91 on page 387.
MVCA_VIRTUALAPPCAPACITYVW	Provides capacity information for a virtual application. See Table 92 on page 387.
MVC_EVENTSVW	Provides event information, such as description and time reported. See Table 73 on page 376.
MVCA_FSRM_VOLUMESUMMARYVW	Provides summary information about File System Viewer volumes. Table 93 on page 388.
MVCA_FSRM_AGESUMMARYVW	Provides summary information about ages in File System Viewer. See Table 94 on page 388.
MVCA_FSRM_EXTDETAILSUMMARYVW	Provides summary information about extent details. See Table 95 on page 389.
MVCA_FSRM_DIRDETAILSUMMARYVW	Provides information about directories in File System Viewer. See Table 96 on page 389.

Table 58 Description of the Report Views (continued)

Materialized View (Tables)	Description
MVCA_FSRM_USERSUMMARYVW	Provides information about users in File System Viewer. See Table 97 on page 389.
MVCA_FSRM_GROUPSUMMARYVW	Provides information about groups in File System Viewer. See Table 98 on page 390.
MVCA_FSRM_TOPNFILES	
MVCA_FSRM_AGEDFILEDETAILS	Provides information about the age properties of files in File System Viewer. See Table 100 on page 391.
MVCA_FSRM_LARGEDIRINFO	
MVC_STORAGEPOOLSUMMARYVW	Provides summary information about storage pools. See Table 64 on page 369.
MVC_STORAGESYSTEMSUMMARYVW	Provides summary information about a storage system. See Table 63 on page 369.
MVC_STORAGEVOLUMESUMMARYVW	Provides summary information about a storage volume. See Table 65 on page 370.
MVC_HOSTDISKDRIVEVW	Provides information about host disk drives. See Table 62 on page 368.
MVC_HOSTVOLUMESUMMARYVW	Provides summary information about host volumes. See Table 61 on page 368.
MVC_DISKEXTENTSUMMARYVW	Provides summary information about disk extents. See Table 81 on page 380.
MVC_STORAGEVOLUMEPORTS	Provides information about storage volume ports. See Table 82 on page 380.
MVC_VOLUMEDISKDRIVEVW	Provides information about volume disk drives. See Table 83 on page 381.
MVC_STORAGEPROCESSORSUMMARYVW	Provides information about storage processors. See Table 84 on page 381.
MVC_DISKDRIVESUMMARYVW	Provides summary information about disk drives. See Table 85 on page 382.
MVC_DISK_EXTENTVW	Provides information about disk extents. See Table 86 on page 383.
MVC_SWITCHSUMMARYVW	Provides summary information about switches. See Table 66 on page 371.

Table 58 Description of the Report Views (continued)

Materialized View (Tables)	Description
MVC_PORTSUMMARYVW	Provides summary information about ports. See Table 67 on page 373.
MVC_ZONESUMMARYVW	Provides summary information about zones. See Table 68 on page 373.
MVC_ZONEVW	Provides information about zones. See Table 69 on page 374.
MVC_PATHVW	Provides path information. See Table 70 on page 375.
MVC_SUBPATHVW	Provides subpath information. See Table 71 on page 375.
MVC_MULTIPATHVW	Provides multipath information. See Table 72 on page 376.
MVC_SWITCHCONFIGVW	Provides switch configuration information. See Table 79 on page 379.
MVC_HOSTCAPACITYVW	Provides host capacity information. See Table 76 on page 378.
MVC_STORAGESYSTEMCONFIGVW	Provides storage system configuration information. See Table 77 on page 378.
MVC_STORAGEPOOLCONFIGVW	Provides information about storage pool configurations. See Table 78 on page 379.
MVIEWCORE_STATUS	Provides information about the core views. The core views are the views starting with mvc_ and mvca. See Table 108 on page 396.
MVIEW_STATUS	Show the status of the materialized views for reports. Table 109 on page 396.

The following tables provide information about each report view:

Table 59 MVC_HOSTSUMMARYVW

Column Name	Type	Description
HOSTID	NUMBER(38)	HostID
HOSTNAME	VARCHAR2(256)	Host Name

Table 59 MVC_HOSTSUMMARYVW (continued)

Column Name	Type	Description
DOMAINID	NUMBER(38)	DomainID
VENDOR	VARCHAR2(256)	Host Vendor
DESCRIPTION	VARCHAR2(1024)	Host Description
STATUS	NUMBER(38)	Operation status (provide map here)
IP	VARCHAR2(16)	Host IP
DNS	VARCHAR2(50)	Host DNS Name
Model	VARCHAR2(256)	Host Model
Version	VARCHAR2(256)	Host Version number
OS	VARCHAR2(24)	Host Operating System
TOTALPHYSICALMEM	NUMBER(38)	Total physical memory
NUMBERPROCESSOR	Number	Number of processors
SUPPORTFLAG	NUMBER(38)	Support flag (unused now)
BASETABLENAME	CHAR(4)	Name of the base table for optional values

Table 60 MVC_CARDSUMMARYVW

Column Name	Type	Description
CardID	NUMBER(38)	CardID
CardName	VARCHAR2(256)	Card Name
ContainerID	NUMBER(38)	Container ID
CardType	VARCHAR2(7)	Card Type (HBA, SCSI)
DomainID	NUMBER(38)	Domain ID
Vendor	VARCHAR2(256)	Card Vendor
Description	VARCHAR2(1024)	Card Description
Status	NUMBER(38)	Operational status
WWN	VARCHAR2(256)	Node WWN
Model	VARCHAR2(256)	Card model

Table 60 MVC_CARDSUMMARYVW (continued)

Column Name	Type	Description
SerialNumber	VARCHAR2(256)	Card Serial Number
Version	VARCHAR2(256)	Card Version
Firmware	VARCHAR2(256)	Firmware version
DriverVersion	VARCHAR2(256)	Driver version
BASETABLENAME	CHAR(4)	Name of the base table for optional values

Table 61 MVC_HOSTVOLUMESUMMARYVW (logical volumes)

Column Name	Type	Description
LogicalVolumeID	NUMBER(38)	Storage Volume ID
LogicalVolumeName	VARCHAR2(256)	Name of the logical volume
DomainID	NUMBER(38)	Domain ID
Description	VARCHAR2(1204)	Description
HostID	NUMBER(38)	Container Host ID
DeviceID	VARCHAR2(254)	Logical Device ID
FileSystemType	VARCHAR2(254)	File System Type
Blocksize	NUMBER(38)	These 3 fields may not be needed for Host
NumberOfBlocks	NUMBER(38)	Logical Volumes
ConsumableBlocks	NUMBER(38)	
BASETABLENAME	CHAR(14)	Name of the base table for optional values

Table 62 MVC_HOSTDISKDRIVEVW

Column Name	Type	Description
HostID	NUMBER(38)	Host ID
DiskDriveID	NUMBER(38)	Disk Drive ID
ExtentID	NUMBER(38)	Disk Partition ID
DiskDrive	VARCHAR2(256)	Disk Drive Name
DiskPartition	VARCHAR2(256)	Disk Partition Name

Table 62 MVC_HOSTDISKDRIVEVW (continued)

Column Name	Type	Description
DiskPartitionDescription	VARCHAR2(1024)	Description of the partition
DiskPartitionSPace	Number	Capacity of the disk partition in megabytes

Table 63 MVC_STORAGESYSTEMSUMMARYVW

Column Name	Type	Description
StorageSystemID	NUMBER(38)	Storage system ID
StorageSystemName	VARCHAR2(256)	Storage system Name
DomainID	NUMBER(38)	Domain ID
Vendor	VARCHAR2(254)	Vendor
Description	VARCHAR2(1024)	Description of the Storage System
Status	NUMBER(38)	Operational Status (provide map here)
IP	VARCHAR2(16)	Not used
Model	VARCHAR2(254)	Model
SerialNumber	VARCHAR2(254)	Serial Number
Version	VARCHAR2(254)	Version
StorageSystemStatus	VARCHAR2(254)	Intrinsic status of the system
ResetCapability	VARCHAR2(254)	Indicating reset capability
ProvisionCapabilities	NUMBER(38)	Provide map here
SupportFlag	NUMBER(38)	Provide map here
BASETABLENAME	Varchar	Name of the base table for optional values

Table 64 MVC_STORGAEPOOLSUMMARYVW

Column Name	Type	Description
StoragePoolID	NUMBER(38)	Storage Pool ID
StoragePoolName	VARCHAR2(256)	Pool Name
StoragePoolDescription	VARCHAR2(1024)	Description of the storage pool

Table 64 MVC_STORGAEPOOLSUMMARYVW (continued)

Column Name	Type	Description
Status	NUMBER(38)	Operational status (provide map)
StorageSystemID	NUMBER(38)	ID of storage system to which pool belongs
PoolSettingID	NUMBER(38)	Storage capabilities ID
ParentPoolID	NUMBER(38)	Parent pool ID, reference to Storage Pool ID
TotalAvailableSpace	NUMBER(38)	Total available space in bytes
CIMPoolID	VARCHAR2(254)	Reserved
PoolType	NUMBER(38)	Pool Type (provide map here)
StorageCapabilityInternalName	VARCHAR2(254)	Internal Name of the capability
NoSinglePTOfFailure	NUMBER(1)	No single point of failure indication
DefaultNoSinglePtOfFailure	NUMBER(1)	Default no single point of failure indication
MinDataRedundancy	NUMBER(18)	Minimum data redundancy indication
MaxDataRedundancy	NUMBER(18)	Maximum data redundancy indication
DefaultDataRedundancy	NUMBER(18)	Default data redundancy indication
MinSpindleRedundancy	NUMBER(18)	
MaxSpindleRedundancy	NUMBER(18)	
DefaultSpindleRedundancy	NUMBER(18)	
MinDeltaReservation	NUMBER(18)	
MaxDeltaReservation	NUMBER(18)	
DefaultDeltaReservation	NUMBER(18)	
StorageCapabilityCommonName	VARCHAR2(256)	Name of the pool capability
StorageCapabilityDescription	VARCHAR2(1024)	Description of the pool capability

Table 65 MVC_STORAGEVOLUMESUMMARYVW

Column Name	Type	Description
StorageVolumeID	NUMBER(38)	StorageVolume ID

Table 65 MVC_STORAGEVOLUMESUMMARYVW (continued)

Column Name	Type	Description
StorageVolumeName	VARCHAR2(256)	StorageVolume Name
DomainId	NUMBER(38)	Domain ID
OID	VARCHAR2(254)	Reserved
Status	NUMBER(38)	Operational status (provide map here)
StorageSystemID	NUMBER(38)	ID of the storage system that contains this volume
StorageCapabilityID	NUMBER(38)	Storage Capability ID
VolumeDeviceID	VARCHAR2(254)	Device ID
AccessType	VARCHAR2(254)	Volume access type
Blocksize	NUMBER(38)	Size per block in bytes
NumberOfBlocks	NUMBER(38)	Total number of blocks in the volume
ConsumableBlocks	NUMBER(38)	Total number of consumable blocks
SeqAccess	NUMBER(1)	Sequential access
Availability	VARCHAR2(254)	Availability indication
StatusInfo	VARCHAR2(254)	Status of the volume
PoolID	NUMBER(38)	ID of the Storage Pool that contains this volume
VolumeType	NUMBER(38)	Type of volume
BASETABLENAME	CHAR(14)	Name of the base table for optional values

Table 66 MVC_SWITCHSUMMARYVW

Column Name	Type	Description
SwitchID	NUMBER(38)	Switch ID
SwitchName	VARCHAR2(256)	Switch Name
DomainID	NUMBER(38)	Domain ID
Vendor	VARCHAR2(254)	Switch Vendor
Description	VARCHAR2(1024)	Description of the Switch
Status	NUMBER(38)	Operational status (provide map here)

Table 66 MVC_SWITCHSUMMARYVW (continued)

Column Name	Type	Description
IP	VARCHAR2(16)	Switch IP
DNS	VARCHAR2(50)	DNS of the Switch
WWN	VARCHAR2(254)	WWN of the Switch
Model	VARCHAR2(254)	Switch Model
SerialNumber	VARCHAR2(254)	Serial Number of the Switch
Version	VARCHAR2(254)	Switch's hardware version
LoginName	VARCHAR2(254)	Login name for this Switch
LoginPwd	VARCHAR2(254)	Login password for this Switch
HardZoningCapability	VARCHAR2(254)	
SoftZoningCapability	VARCHAR2(254)	
ZoningInstalled	NUMBER(1)	
MaxModuleNumber	NUMBER(38)	
CurrentZoningEnforcement	VARCHAR2(254)	
SwitchDomainID	NUMBER(38)	
SwitchStatus	VARCHAR2(254)	
SwitchState	VARCHAR2(254)	
IPGateway	VARCHAR2(254)	
IPNetwork	VARCHAR2(16)	
FCAddress	VARCHAR2(254)	
FCNetmask	VARCHAR2(16)	
SwitchRole	VARCHAR2(254)	
ProvisionSupportFlag	NUMBER(1)	
FabricWWN	Varchar	
FabricID	NUMBER(38)	Fabric ID
BASETABLENAME	CHAR(6)	Name of the base table for optional values

Table 67 MVC_PORTSUMMARYVW

Column Name	Type	Description
PortID	NUMBER(38)	Port ID
PortName	VARCHAR2(256)	Port Name
DomainID	NUMBER(38)	Domain ID
Description	VARCHAR2(1024)	Description
Status	NUMBER(38)	Operational status (provide map)
WWN	VARCHAR2(32)	Port WWN
ContainerID	NUMBER(38)	Container ID of this port
TargetPort	NUMBER(38)	Target Port ID (Never populated)
Connected_To_WWN	VARCHAR2(32)	WWN of connected port
Device_ID	VARCHAR2(254)	Port Device ID
PortState	VARCHAR2(254)	
PortStatus	VARCHAR2(254)	
Physical_State	VARCHAR2(254)	
Port_ID	NUMBER(38)	
Port_Speed	NUMBER(38)	
Max_Speed	NUMBER(38)	Port Max Speed, bit/second
PortNumber	NUMBER(18)	
SCSIPort	NUMBER(18)	
ConnectedToNodeWWN	VARCHAR2(32)	
PortType	VARCHAR2(254)	
BASETABLENAME	CHAR(4)	Name of the base table for optional values

Table 68 MVC_ZONESUMMARYVW

Column Name	Type	Description
ZoneID	NUMBER(38)	Zone ID

Table 68 MVC_ZONESUMMARYVW (continued)

Column Name	Type	Description
ZoneName	VARCHAR2(254)	Zone Name
DominalID	NUMBER(38)	Domain ID (currently only one domain)
CimClassName	VARCHAR2(28)	
Status	NUMBER(38)	ApplQ status
ActiveZone	VARCHAR2(3)	
ZoneType	VARCHAR2(254)	
ProtocolType	VARCHAR2(254)	
ReadOnly	NUMBER(1)	
FabricID	NUMBER(38)	Fabric ID
FabricWWN	VARCHAR2(254)	Fabric WWN
FabricCName	VARCHAR2(256)	
ZoneCapID	NUMBER(38)	Zone capability ID
ZoneCapabilitiesName	VARCHAR2(254)	Name of the zone capabilities
ZC_MaxName_length	NUMBER(18)	Name length limit
MaxZoneSets	NUMBER(18)	Number of maximum zone sets
MaxZones	NUMBER(18)	
MaxZoneMembers	NUMBER(18)	
MaxZonePerZoneSet	NUMBER(18)	
MaxZoneAliases	NUMBER(18)	
EnhencedZoning	NUMBER(1)	

Table 69 MVC_ZONEVW

Column Name	Type	Description
ZoneSetID	NUMBER(38)	Zone set ID
ZoneSetName	VARCHAR2(256)	Zone set Name
ZoneID	NUMBER(38)	Zone ID
FabricID	NUMBER(38)	ID of the fabric which the zone belongs

Table 69 MVC_ZONEVW (continued)

Column Name	Type	Description
ZoneMemberID	NUMBER(38)	Zone member ID
ZoneMemberName	VARCHAR2(254)	Name of the zone member
ZoneMemberType	VARCHAR2(254)	Type of the zone member
ZoneMemberInFabric	NUMBER(1)	
ZonePortWWN	VARCHAR2(32)	WWN of the zone port
ZoneAlias	VARCHAR2(256)	

Table 70 MVC_PATHVW

Column Name	Type	Description
PathID	NUMBER(38)	Path ID
HostID	NUMBER(38)	Host ID on this path
LogicalVolumeID	NUMBER(38)	Logical Volume ID
DiskPartitionID	NUMBER(38)	Disk Partition ID if raw partition is present
IsMounted	VARCHAR2(5)	(TRUE, FALSE)
ContainerPathID	NUMBER(38)	Currently not in use
ApplicationID	NUMBER(38)	Application ID for the application file
ApplicationFileID	NUMBER(38)	Application File ID
FilePath	VARCHAR2(256)	Application file path
FileName	VARCHAR2(128)	Application file name

Table 71 MVC_SUBPATHVW

Column Name	Type	Description
SubPathID	NUMBER(38)	SubPath ID
PathID	NUMBER(38)	Parent Path ID
HostID	NUMBER(38)	Host ID
DiskDriveID	NUMBER(38)	Disk Drive ID

Table 71 MVC_SUBPATHVW (continued)

Column Name	Type	Description
HBACardID	NUMBER(38)	HBA Card ID
HBAPortID	NUMBER(38)	HBA Port ID
HostSwitchPortID	NUMBER(38)	ID of Host Switch Port
SystemSwitchPortID	NUMBER(38)	ID of the system switch port
StorageSystemPortID	NUMBER(38)	
StorageVolumeID	NUMBER(38)	
StorageSystemID	NUMBER(38)	
LUN	NUMBER(38)	
IsLocal	VARCHAR2(6)	
FabricID	NUMBER(38)	
MultipathDeviceID	NUMBER(38)	
PathSwitchID	NUMBER(38)	

Table 72 MVC_MULTIPATHVW

Column Name	Type	Description
PathID	NUMBER(38)	Path ID
MultiPathDiskExtentID	NUMBER(38)	ID of multipath disk extent
VolumeManagerVolumeID	NUMBER(38)	Volume Manager Volume ID
VxvmDiskPartitionID	NUMBER(38)	ID of volume manager volume disk partition
multipathDeviceID	NUMBER(38)	

Table 73 MVC_EVENTSVW

Column Name	Type	Description
DomainID	NUMBER(38)	
ID	NUMBER(38)	
Element_ID	NUMBER(38)	
Device_ID	NUMBER(38)	

Table 73 MVC_EVENTS VW (continued)

Column Name	Type	Description
Summary_Text	VARCHAR2(4000)	
Time_Reported	Date	
Severity	NUMBER(38)	
Cleared	NUMBER(1)	
Source	VARCHAR2(254)	
Type	NUMBER(38)	
SubType	NUMBER(38)	
Probable_Cause_Description	VARCHAR2(4000)	
Recommanded_Actions	VARCHAR2(254)	
Correlated_Indications	VARCHAR2(254)	
Type_Description	VARCHAR2(254)	
Element_type	NUMBER(38)	
Name	VARCHAR2(256)	
Vendor	VARCHAR2(254)	
Model	VARCHAR2(254)	

Table 74 MVC_ORGANIZATION VW

Column Name	Type	Description
ORGID	NUMBER(38)	ID from Organization table
ORGNAME	VARCHAR2(256)	Organization name
ORGDESCRIPTION	VARCHAR2(4000)	Description of this Organization
ELEMENTID	NUMBER(38)	Element ID mapped to this Organization
DOMAINID	NUMBER(38)	Domain ID

Table 75 MVC_ORGRELATION VW

Column Name	Type	Description
ParentOrgID	NUMBER(38)	Parent Organization ID

Table 75 MVC_ORGRELATIONVW (continued)

Column Name	Type	Description
ChildOrgID	NUMBER(38)	Child Organization ID
DOMAINID	NUMBER(38)	Domain ID

Table 76 MVC_HOSTCAPACITYVW

Column Name	Type	Description
HostID	NUMBER(38)	Host ID
VolumeID	NUMBER(38)	Logical Volume ID
TimeStamp	TIMESTAMP	Time of data collection
Total	NUMBER(38)	Total capacity in megabytes
Used	NUMBER(38)	Used capacity in megabytes
Free	NUMBER(38)	Free capacity in megabytes

Table 77 MVC_STORAGESYSTEMCONFIGVW

Column Name	Type	Description
StorageSystemID	NUMBER(38)	Storage System ID
CollectionTime	TIMESTAMP(6)	Configuration statistics collection time
MaskedMB	NUMBER	Masked storage in megabytes
ExportedMB	NUMBER	Storage exposed in megabytes
UnexportedMB	NUMBER	Volume unexposed in megabytes
AvailableMB	NUMBER	Available storage left on storage pool
ProvisionedMB	NUMBER	Provisioned storage in megabytes
RawStorageMB	NUMBER	Unconfigured storage in megabytes
TotalMB	NUMBER	Total storage in megabytes
MainframeStorageMB	NUMBER	Mainframe storage in megabytes
AvailablePorts	NUMBER(38)	Number of available ports
ConnectedPorts	NUMBER(38)	Number of used ports

Table 77 MVC_STORAGESYSTEMCONFIGVW (continued)

Column Name	Type	Description
TotalPorts	NUMBER(38)	Total system ports

Table 78 MVC_STORAGEPOOLCONFIGVW

Column Name	Type	Description
StoragePoolID	NUMBER(38)	Storage Pool ID
CollectionTime	TIMESTAMP (6)	Configuration stats collection time
MaskedMB	NUMBER	Masked Storage in megabytes
ExportedMB	NUMBER	Storage exposed in megabytes
UnexportedMB	NUMBER	Volume unexported in megabytes
AvailableMB	NUMBER	Available storage left on storage pool
ProvisionedMB	NUMBER	Provisioned storage in megabytes
TotalMB	NUMBER	Total storage in megabytes
StorageCapability	VARCHAR2(255)	Storage pool capability
MainframePool	VARCHAR2(13)	Indicate whether reserved for mainframe: MAINFRAMEPOOL

Table 79 MVC_SWITCHCONFIGVW

Column Name	Type	Description
SwitchID	NUMBER(38)	Switch ID
AvailablePorts	NUMBER(38)	Number of available ports
ConnectedPorts	NUMBER(38)	Number of used ports
TotalPorts	NUMBER(38)	Number of total ports of this switch

Table 80 MVC_OPTIONALTABLEVW

Column Name	Type	Description
BaseTableName	VARCHAR2(30)	Name of the base table
BaseTableID	NUMBER(38)	ID of the base table record

Table 80 MVC_OPTIONALTABLEVW (continued)

Column Name	Type	Description
OptionalName	VARCHAR2(50)	Name of the optional value
OptionalValue	VARCHAR2(4000)	Value of the optional value

Table 81 MVC_DISKEXTENTSUMMARYVW

Column Name	Type	Description
DiskExtentID	NUMBER(38)	Disk Extent ID
DiskEnxtentName	VARCHAR2(256)	Name of the extent
DomainID	NUMBER(38)	Domain ID
DiskExtentDescription	VARCHAR2(1024)	Description of the extent
Status	NUMBER(38)	Operational status
Access_Type	VARCHAR2(254)	Access type
BlockSize	NUMBER(38)	Block Size
Number_Of_Blocks	NUMBER(38)	Number of total blocks
Consumable_Blocks	NUMBER(38)	Number of usable blocks
StoragePoolID	NUMBER(38)	Storage Pool ID
SystemID	NUMBER(38)	Container System ID
BASETABLENAME	CHAR(11)	Name of the base table

Table 82 MVC_STORAGEVOLUMEPORTS

Column Name	Type	Description
ID	NUMBER(38)	
Storage_Volume_ID	NUMBER(38)	Storage Volume ID
Port_ID	NUMBER(38)	
LUN_ID	NUMBER(38)	
Access_Mode	Varchar(254)	
Access_State	Varchar(254)	
Host_Group	Varchar(254)	

Table 82 MVC_STORAGEVOLUMEPORTS (continued)

Column Name	Type	Description
Host_Group_Name	Varchar(254)	
Host_Group_Modes	Varchar(1024)	

Table 83 MVC_VOLUMEDISKDRIVEVW

Column Name	Type	Description
VolumeID	NUMBER(38)	Storage Volume ID
DiskDriveID	NUMBER(38)	Disk Drive ID
ExtentID	NUMBER(38)	Disk Extent ID

Table 84 MVC_STORAGEPROCESSORSUMMARYVW

Column Name	Type	Description
SystemProcessorID	NUMBER(38)	Storage System Processor ID
SystemProcessorName	VARCHAR2(256)	Name of the system processor
DomainID	NUMBER(38)	Domain ID
Vendor	VARCHAR2(254)	Vendor
Description	VARCHAR2(1024)	
Status	NUMBER(38)	Operational status
IP	VARCHAR2(16)	
DNS	VARCHAR2	
WWN	VARCHAR(16)	
Model	VARCHAR2(254)	
PowerManagement	VARCHAR2(254)	
SerialNumber	VARCHAR2(254)	
Version	VARCHAR2(254)	
ContainerID	NUMBER(38)	Container system ID
ProcessorStatus	NUMBER	Status of the processor

Table 84 MVC_STORAGEPROCESSORSUMMARYVW (continued)

Column Name	Type	Description
ResetCapability	VARCHAR2(254)	Reset Capability
Roles	VARCHAR2(254)	Roles
ProviderTag	VARCHAR2(254)	Provider name tag
SupportFlags	NUMBER(38)	Support flags
BASETABLENAME	CHAR(14)	Name of the base table

Table 85 MVC_DISKDRIVESUMMARYVW (sheet 1 of 2)

Column Name	Type	Description
DiskDriveID	NUMBER(38)	Disk Drive ID
DiskDriveName	VARCHAR2(256)	Name of the Disk Drive
DomainID	NUMBER(38)	Domain ID
OID	VARCHAR2(254)	Object ID of the Disk Drive
Vendor	VARCHAR2(254)	Vendor
Description	VARCHAR2(1024)	Description
Status	NUMBER(38)	Operational status
Model	VARCHAR2(254)	Disk Drive Model
Name	VARCHAR2(254)	Name coming from disk drive
CardID	NUMBER(38)	Card ID
DiskDriveStatus	VARCHAR2(254)	Disk Drive Status
SCSIBus	NUMBER(38)	SCSI Bus
SCSILUN	NUMBER(38)	SCSI LUN
SCSITargetID	NUMBER(38)	SCSI target ID
SCSIPort	NUMBER(38)	SCSI PORT ID
SystemID	NUMBER(38)	Container System ID
MaxMediaSize	NUMBER(38)	Maximum media size
MaxBlockSize	NUMBER(38)	Maximum block size
MinBlockSize	NUMBER(38)	Minimum block size

Table 85 MVC_DISKDRIVESUMMARYVW (continued) (sheet 2 of 2)

Column Name	Type	Description
EnableStatus	VARCHAR2(254)	
Availability	VARCHAR2(254)	
BASETABLENAME	CHAR(11)	Name of the base table

Table 86 MVC_DISK_EXTENTVW

Column Name	Type	Description
ExtentID	NUMBER	Disk Extent ID
ContainerExtentID	NUMBER	Container Extent ID
DiskID	NUMBER	Disk Drive ID

Table 87 MVC_ASSETSUMMARY

Column Name	Type	Description
DOMAINID	NUMBER(38)	Domain ID
ASSETID	NUMBER(38)	ID of the Asset depending on Assetclass for example if the assetclass is Host then this is host.id
ASSETCLASS	VARCHAR2(13)	The Asset class are "HOST", "APPLICATION", "STORAGESYSTEM" or "SWITCH"
NAME	VARCHAR2(256)	Name of the Host for example in case of Host this host name.
DATECREATED	DATE	Creation Date
DATELASTMODIFIED	DATE	Date last modified.
DESCRIPTION	VARCHAR2(255)	Asset description
STATUS	VARCHAR2(8)	Asset status. Can be 'NEW', 'MISSING', 'REPAIRED', 'IN USE'
VENDOR	VARCHAR2(254)	Asset Vendor
MODEL	VARCHAR2(254)	Asset Model
SERIALNUMBER	VARCHAR2(254)	Asset serial number

Table 87 MVC_ASSETSUMMARY (continued)

Column Name	Type	Description
BARCODE	VARCHAR2(255)	Asset bar code
ASSETCODE	VARCHAR2(255)	Asset code
ASSETTYPE	VARCHAR2(255)	Asset type
ASSETTAG	VARCHAR2(255)	Asset Tag
ASSETCATEGORY	VARCHAR2(255)	Asset Category
GEOGRAPHICLOCATION	VARCHAR2(255)	Asset location
STORAGETIERCLASSIFICATION	VARCHAR2(255)	Asset storage Tier Name for example "Ultra High Availability"
STORAGETIERCOSTPERGB	NUMBER(36,2)	Asset storage Tier cost
DEPARTMENTNO	VARCHAR2(255)	Asset department no
DEPARTMENTNAME	VARCHAR2(255)	Asset department Name
PERCENTAGEOWNED	NUMBER(5, 2)	Percentage owned by Department
ADMINISTRATOR	VARCHAR2(255)	Asset maintained by
STAFFNAME	VARCHAR2(255)	Staff Name
STAFFPHONENUMBER	VARCHAR2(255)	Staff PH#
STAFFDEPARTMENT	VARCHAR2(255)	Staff Department
STAFFEMAIL	VARCHAR2(255)	Staff E-mail
RACKNUMBER	VARCHAR2(100)	
FLOOR	VARCHAR2(100)	
DATACENTER	VARCHAR2(100)	
CITY	VARCHAR2(100)	
REGION	VARCHAR2(100)	
COUNTRY	VARCHAR2(50)	
CONTINENT	VARCHAR2(20)	
ADDRESS	VARCHAR2(1024)	
ZIPCODE	VARCHAR2(16)	
LICENSE	VARCHAR2(4000)	

Table 87 MVC_ASSETSUMMARY (continued)

Column Name	Type	Description
PURCHASEORDERNUMBER	VARCHAR2(255)	Asset Purchase Order no
DATEPURCHASED	DATE	Asset Date purchased
COST	NUMBER(36,2)	Asset cost
VALUE	NUMBER(36,2)	
SALVAGECOST	NUMBER(36,2)	Asset Deprecated Salvage Cost
DEPRECITIONPERIOD	NUMBER(18)	Asset Deprecation Period
DEPRECITIONMETHOD	NUMBER(16)	Asset Deprecation Method ('Straight line', 'Fixed declining', 'Double declining')
DEPRECITEDVALUE	NUMBER(36,2)	Asset Depreciated Value
RESELLER	VARCHAR2(255)	Asset Reseller
COMMENTS	VARCHAR2(4000)	Comments
ASSETFIXCOSTTAXPERDEPTPE RYEAR	NUMBER	Asset Fixed cost tax per department per Year
STORAGFIXCOSTTAXPERDEP TPERYEAR	NUMBER	Storage Fixed cost tax per department per Year

Table 88 MVC_APPLICATIONSUMMARYVW

Column Name	Type	Description
APPLICATIONID	NUMBER(38)	Application ID
APPLICATIONNAME	VARCHAR2(256)	Application Name
DOMAINID	NUMBER(38)	Domain ID
APPLICATIONTYPE	NUMBER(38)	Application type
VENDOR	VARCHAR2(254)	Vendor
DESCRIPTION	VARCHAR2(1024)	Description of the application
STATUS	NUMBER(38)	Operational status
BUILDNUMBER	VARCHAR2(254)	Software build number
SERIALNUMBER	VARCHAR2(254)	Software serial number

Table 88 MVC_APPLICATIONSUMMARYVW (continued)

Column Name	Type	Description
TARGETOS	VARCHAR2(254)	Target operating system
PRODUCTNAME	VARCHAR2(254)	Product name
VERSION	VARCHAR2(254)	Software version
IDENTIFICATIONCODE	VARCHAR2(254)	Software identification code
SOFTWAREELEMENTID	VARCHAR2(254)	Software element ID
HOSTID	NUMBER(38)	ID of host where the application is running from

Table 89 MVC_UNITACCESSVW

Column Name	Type	Description
ID	NUMBER(38)	
STORAGE_VOLUME_ID	NUMBER(38)	
STORAGE_SYSTEM_PORT_ID	NUMBER(38)	
HBA_PORT_ID	NUMBER(38)	
ACCESS_MODE	VARCHAR(254)	
INITIATOR	VARCHAR(254)	
INITIATOR_FORMAT	NUMBER(38)	
UNIT_NUMBER	NUMBER(38)	
HOST_GROUP	VARCHAR(254)	
HOST_GROUP_NAME	VARCHAR(254)	
HOST_GROUP_MODES	VARCHAR(1024)	

Application Core Views

Table 90 MVCA_DBAPPCAPACITYVW

Column Name	Type	Description
DBApplicationID	NUMBER(38)	
HostID	NUMBER(38)	

Table 90 MVCA_DBAPPCAPACITYVW (continued)

Column Name	Type	Description
CapacityType	Varchar2(7)	Type of capacity data
Timestamp	Timestamp(10)	
DBInstanceName	Varchar2(255)	
DBLogicalName	Varchar2(255)	
DBPhysicalName	Varchar2(512)	
TotalMB	NUMBER	
FreeMB	NUMBER	
UsedMB	NUMBER	

Table 91 MVCA_EXCHAPPCAPACITYVW

Column Name	Type	Description
ExchangeAppID	NUMBER(38)	
HostID	NUMBER(38)	
CapacityType	Varchar2(7)	
Timestamp	Timestamp	
StorageGroupID	NUMBER(38)	
ExchangeFilePath	Varchar(512)	
TotalMB	NUMBER	
FreeMB	NUMBER	
UsedMB	NUMBER	

Table 92 MVCA_VIRTUALAPPCAPACITYVW

Column Name	Type	Description
VirtualAppID	NUMBER(38)	
HostID	NUMBER(38)	

Table 92 MVCA_VIRTUALAPPCAPACITYVW (continued)

Column Name	Type	Description
Timestamp	Date	
VirtualPath	Varchar2(512)	
TotalMB	NUMBER	
FreeMB	NUMBER	
UsedMB	NUMBER	

Table 93 MVCA_FSRM_VOLUMESUMMARYVW

Column Name	Type	Description
VolumeID	NUMBER(38)	
VolumeName	Varchar2(256)	
FSID	NUMBER(38)	
TotalDirectories	NUMBER(38)	
TotalFiles	NUMBER(38)	
DomainID	NUMBER(38)	
Timestamp	Timestamp(6)	

Table 94 MVCA_FSRM_AGESUMMARYVW

Column Name	Type	Description
AgeID	NUMBER(38)	
VolumeID	NUMBER(38)	
TotalFiles	NUMBER(38)	
TotalSize	NUMBER(38)	
Timestamp	Timestamp(6)	

Table 95 MVCA_FSRM_EXTDETAILSUMMARYVW

Column Name	Type	Description
ExtName	Varchar2(254)	
VolumeID	NUMBER(38)	
TotalFiles	NUMBER(38)	
TotalSize	NUMBER(38)	
Timestamp	Timestamp(6)	

Table 96 MVCA_FSRM_DIRDETAILSUMMARYVW

Column Name	Type	Description
DirKey	NUMBER(38)	
ParentKey	NUMBER(38)	
DirName	Varchar2(254)	
DirLSevel	NUMBER(38)	
DirSize	NUMBER(38)	
TotalSubDirectories	NUMBER(38)	
TotalFiles	NUMBER(38)	
VolumeID	NUMBER(38)	
Timestamp	Timestamp(6)	

Table 97 MVCA_FSRM_USERSUMMARYVW

Column Name	Type	Description
UserID	NUMBER(38)	
FSID	NUMBER(38)	
UserProvidederID	Varchar2(254)	
UserName	Varchar2(254)	
DirName	Varchar2(254)	

Table 97 MVCA_FSRM_USERSUMMARYVW (continued)

Column Name	Type	Description
Department	Varchar2(254)	
Email	Varchar2(254)	
Quota	NUMBER(38)	
DomainID	NUMBER(38)	

Table 98 MVCA_FSRM_GROUPSUMMARYVW

Column Name	Type	Description
GroupID	NUMBER(38)	
FSID	NUMBER(38)	
GroupName	Varchar2(254)	
DirName	Varchar2(254)	
Contact	Varchar2(254)	
Department	Varchar2(254)	
Email	Varchar2(254)	
Quota	NUMBER(38)	
DomainID	NUMBER(38)	

Table 99 MVCA_FSRM_TOPNFILES

Column Name	Type	Description
GroupID	NUMBER(38)	
FSID	NUMBER(38)	
GroupName	Varchar2(254)	
DirName	Varchar2(254)	
Contact	Varchar2(254)	
Department	Varchar2(254)	
Email	Varchar2(254)	
Quota	NUMBER(38)	

Table 99 MVCA_FSRM_TOPNFILES (continued)

Column Name	Type	Description
DomainID	NUMBER(38)	

Table 100 MVCA_FSRM_AGEDFILEDETAILS

Column Name	Type	Description
VolumeID	NUMBER(38)	
FileName	Varchar2(254)	
FileSize	NUMBER(38)	
FileAge	NUMBER(38)	
Timestamp	Timestamp(6)	
DomainID	NUMBER(38)	

Table 101 MVCA_FSRM_LARGEDIRINFO

Column Name	Type	Description
VolumeID	NUMBER(38)	
DirName	Varchar2(256)	
DirSize	NUMBER(38)	
TotalFiles	NUMBER(38)	
TotalDirs	NUMBER(38)	
Timestamp	Timestamp(6)	
DomainID	NUMBER(38)	

Table 102 MVCA_BU_MASTERSERVERSUMMARY

Column Name	Type	Description
MasterServerID	NUMBER(38)	
MasterServerName	Varchar2(256)	
HostID	NUMBER(38)	

Table 102 MVCA_BU_MASTERSERVERSUMMARY (continued)

Vendor	Varchar2(254)	
Description	Varchar2(1024)	
Status	NUMBER	
ProductName	Varchar2(254)	
LicenseKey	Varchar2(256)	
LicenseFeatures	Varchar2(256)	
DomainID	NUMBER(38)	

Table 103 MVCA_BU_MEDIASERVERSUMMARY

Column Name	Type	Description
MediaServerID	NUMBER	
MediaServerName	Varchar2(256)	
MasterServerID	NUMBER	
HostID	NUMBER(38)	
Vendor	Varchar2(254)	
Description	Varchar2(1024)	
Status	NUMBER	
ProductName	Varchar2(254)	
LicenseKey	Varchar2(256)	
LicenseFeatures	Varchar2(256)	
DomainID	NUMBER(38)	

Table 104 MVCA_BU_CLIENTSUMMARY

Column Name	Type	Description
ClientID	NUMBER(38)	
ClientName	Varchar2(256)	
MasterServerID	NUMBER	
HostID	NUMBER(38)	

Table 104 MVCA_BU_CLIENTSUMMARY (continued)

Vendor	Varchar2(254)	
Description	Varchar2(1024)	
Status	NUMBER	
ProductName	Varchar2(254)	
DomainID	NUMBER(38)	

Table 105 MVCA_BU_MEDIASUMMARY

Column Name	Type	Description
MediaID	NUMBER(38)	
MediaName	Varchar2(256)	
TapeLibraryID	NUMBER	
PoolID	NUMBER	
MasterServerID	NUMBER(38)	
TLMediaID	Varchar2(32)	
Type	Varchar2(32)	
Barcode	Varchar2(32)	
MediaPoolName	Varchar2(256)	
RobotType	Varchar2(32)	
RobotNumber	NUMBER	
RobotSlot	NUMBER	
RobotHost	Varchar2(128)	
VolumeGroup	Varchar2(64)	
Created	Date	
Assigned	Date	
LastMounted	Date	
FisrtMounted	Date	
ExpirationDate	Date	
NumberOfMounths	NUMBER	

Table 105 MVCA_BU_MEDIASUMMARY (continued)

Column Name	Type	Description
MaxMountsAllocated	NUMBER	
Density	Varchar2(64)	
TimeAllocated	Date	
LastWritten	Date	
Expir	Varchar2(32)	
LastRead	Date	
Mbytes	NUMBER	
NIImages	NUMBER	
VIImages	NUMBER	
RL	Varchar2(64)	
TotalRestores	NUMBER	
MediaStatus	Varchar2(16)	
Vendor	Varchar2(254)	
Description	Varchar2(1024)	
DomainID	NUMBER(38)	

Table 106 MVCA_BU_JOBSUMMARY

Column Name	Type	Description
JobID	NUMBER(38)	
TemplateID	NUMBER	
TemplateName	Varchar2(64)	
MasterServerID	NUMBER	
ClientID	NUMBER	
BUJobID	NUMBER	
JobState	Varchar2(16)	
JobStatus	Varchar2(16)	
ScheduleName	Varchar2(32)	

Table 106 MVCA_BU_JOBSUMMARY (continued)

Column Name	Type	Description
StorageUnit	Varchar2(64)	
BUTargetServer	Varchar2(128)	
FilesLastWritten	NUMBER	
StartTime	Date	
EndTime	Date	
Description	Varchar2(256)	
Time	NUMBER	
RetentionPeriod	Varchar2(16)	
Compression	Varchar2(16)	
Priority	Varchar2(16)	
KBLastWritten	NUMBER	
FileListCount	NUMBER	

Table 107 MVCA_BU_LIBRARYSUMMARY

Column Name	Type	Description
TapeLibraryID	NUMBER	
TapeLibraryName	Varchar2(256)	
Vendor	Varchar2(254)	
Description	Varchar2(1024)	
MediaServerID	NUMBER	
MasterServerID	NUMBER	
Type	Varchar2(64)	
RobotType	Varchar2(64)	
RobotNumber	NUMBER	
TotalNoOfSlots	NUMBER	
TotalSlotsInUse	NUMBER	
TotalNumberOfDrives	NUMBER	

Table 107 MVCA_BU_LIBRARYSUMMARY (continued)

Column Name	Type	Description
RobotDevicePath	Varchar2(128)	
DomainID	NUMBER(38)	

Table 108 MVIEWCORE_STATUS

Name	Type
MVIEWNAME	NOT NULL VARCHAR2(30)
LAST_REFRESH_TIME	DATE
TOTALREFRESHTIME	VARCHAR2(32)
STATUS	VARCHAR2(10)

Table 109 MVIEW_STATUS

Name	Type
MVIEWNAME	NOT NULL VARCHAR2(30)
LAST_REFRESH_TIME	DATE
TOTALREFRESHTIME	VARCHAR2(32)
STATUS	VARCHAR2(10)

Views from Previous Releases

In this release, the materialized views were renamed, revised and in some cases removed. The following views were dropped from this release:

- MV_STORAGESYSTEMCAPSUMMARYVW
- MV_HOSTDETAILVW
- MV_UNITACCESSVW

The following table provides the views from earlier releases. Although you can still use these views for your reports created in previous versions, it strongly recommended you use the new views for any new report development. See the following tables for the corresponding new views.

Table 110 Views from Previous Releases

View Name	Corresponding New View
MV_ASSETCOUNTVW	No replacement
MV_ASSETSUMMARYVW	MVC_ASSETSUMMARYVW
MV_ASSETSUMMARYVW	MVC_ASSETSUMMARYVW
MV_HOSTSUMMARYVW , MV_HOSTCONNECTIVITYV W, MV_HOSTDETAILVW	MVC_HOSTSUMMARYVW
MV_APPLICATIONVW	Information about applications is now in spread across several materialized views.
MV_EVENTVW	MVC_EVENTS VW
MV_FILESERVERVW	Information about file servers is now spread across several materialized views.
MV_HOSTSTORAGEALLOC ATIONVW	Information spread across several materialized views, such as MVC_HOSTSUMMARYVW.
MV_HOSTSTORAGESUMM ARYVW	Information about storage on hosts is now spread across several reports.
MV_HOSTHOSTSTORAGEO CATIONVW	No replacement. Information in several views now.
MV_HOSTSTORAGESUMM ARYVW	MVC_HOSTSUMMARYVW
MV_STORAGESYSTEMPORT UTILVW	No replacement. Information in several views, such as MVC_STORAGESYSTEMCONFIGVW.
MV_SWITCHDETAILVW	MVC_SWITCHSUMMARYVW
MV_ASSETSUMMARYVW	MVC_ASSETSUMMARY
MV_STORAGESYSTEMCAPS UMMARYVW	STORAGEVOLUMESUMMARYVW

13 Viewing Performance Data

IMPORTANT: Depending on your license, Performance Manager may not be available. See the “List of Features” to determine if you have access to Performance Manager. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).

This chapter describes the following:

- “[About Performance Manager](#)” on page 399
- “[Accessing Performance Manager](#)” on page 400
- “[Creating Performance Charts](#)” on page 400
- “[About the Toolbars in Performance Manager](#)” on page 401
- “[Comparing the Performance of Different Elements](#)” on page 404
- “[Viewing Summary Charts](#)” on page 404
- “[Viewing Trending Information for Performance](#)” on page 405
- “[Setting Thresholds in Performance Charts](#)” on page 406
- “[Removing Performance Data from a Graph](#)” on page 406
- “[Setting a Custom Period](#)” on page 407
- “[About the Monitoring Options](#)” on page 408
- “[Task Already Running on Element](#)” on page 417
- “[Monitoring with Direct Attached Storage](#)” on page 418
- “[Supported Host Configurations for Monitoring](#)” on page 418
- “[Sudden Dips Displayed in Certain Charts in Performance Manager](#)” on page 420
- “[Values Continue to Increase in Charts for Aggregated Drives and Aggregate Volumes](#)” on page 421

About Performance Manager

IMPORTANT: Depending on your license, Performance Manager may not be available. See the “List of Features” to determine if you have access to Performance Manager. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center** in HP Storage Essentials).

Performance Manager provides a graphical representation of the performance history of a managed element, such as bytes transmitted per second for a switch.

You can manipulate the charts, so they show a different reporting period and frequency. For example, you could show the performance of an element over the past 24 hours with an hourly monitoring frequency.

Keep in mind the following:

- If you see “There is not enough data to produce a chart [chart_title] at this time.” message, lessen the frequency option, so that the amount of time listed in the **Frequency** drop-down menu has passed before you view the chart again. For example, if the **Frequency** menu displays hourly, you need to wait an hour for data to appear in the chart. If the frequency option is **Default**, the information is displayed according to the interval of the performance collector that gathers information for the graph. For example, if you set the collector for the graph to every 15 minutes and the period set to **Last Hour**, the default setting displays four data points in the graph. See “[Managing Performance Collection](#)” on page 136 for more information about setting the performance collectors.
- Verify the performance collector for that element is enabled. See “[Managing Performance Collection](#)” on page 136 for more information about enabling performance collectors.
- If you select the **Default** option from the Frequency combo box in Performance Manager and the maximum number of data points is more than 20, the management server displays statistics averaged over a three-minute interval. All data points collected within the three-minute period are averaged together to display one data point. For example, assume you are viewing the performance of the read bandwidth on a storage system and you selected the **Default** option from the Frequency combo box in Performance Manager. If there are more than 20 data points within the selected time period, Performance Manager employs an averaging algorithm to return an averaged representation of data in a three-minute period.
- Direct attached storage ports are not displayed in the storage tree in the bottom left pane.
- Performance Manager is not available to file servers.
- To learn more about the buttons in the toolbar, see “[About the Toolbars in Performance Manager](#)” on page 401.
- The Aggregated volume and aggregated drive filters are no longer shown for Engenio 5884 storage systems or any of its controllers. These filters were available in builds earlier than 4.0 of the management server.

IMPORTANT: All collectors are stopped during Discovery Data Collection. This means that during Discovery Data Collection data for Performance Manager is not updated. Historical collectors, such as those available from the Configuration tab, are restarted when they are stopped during Discovery Data Collection. Any charts that were active in Performance Manager when Discovery Data Collection was started are not restarted.


Accessing Performance Manager

To access Performance Manager, click **Optimize > Storage Essentials > Performance Manager**.

Creating Performance Charts

To create a performance chart for an element:

1. Access Performance Manager as described in “[Accessing Performance Manager](#)” on page 400.
2. Select the element you want to monitor.

3. Under the Monitoring tab in the lower-left pane, select the element again. In some instances, you may need to select an element's port, such as a switch.
4. In the pane under the tree, select a monitoring option.
See "[About the Monitoring Options](#)" on page 408 for more information.
5. Use the Chart, Unit and Frequency combo box to modify the chart. When you are done with your selections, click the  button in the lower pane. To learn more about these features, see "[About the Toolbars in Performance Manager](#)" on page 401
6. To monitor more than one element in a chart, see "[Comparing the Performance of Different Elements](#)" on page 404.

About the Toolbars in Performance Manager

Performance Manager provides two toolbars. One in the upper pane and another in the lower pane. The toolbar in the upper pane is the same as the one in System Manager. See "[About the Toolbar in System Manager](#)" on page 164 for information about the toolbar in the upper pane.

The toolbar in the lower pane provides the following information:

Table 111 Toolbar in Lower Pane of Performance Manager






Icon	Description
	Saves the graph in the cache. If you close the Web browser window, the graph is no longer saved.
	Lets you print a graph.
	Clears the graph of the elements you have selected.
	<p>Lets you determine the type of graph displayed. Select one of the following options and then click the  button for your selection to take effect:</p> <ul style="list-style-type: none"> • Bar - Displays each data point as a bar. The data for the different elements is displayed side by side. • Line - Displays each data point as a dot with a line connected to the previous data points. The data for the different elements for a specific point in time is displayed in the same column. • Summary - Displays a single line that summarizes the values for a single statistic. Multiple statistics can be shown with multiple lines. See "Viewing Summary Charts" on page 404.

Table 111 Toolbar in Lower Pane of Performance Manager (continued)

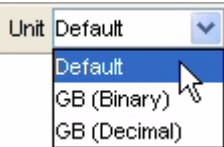

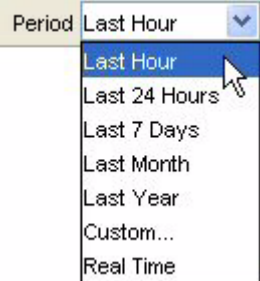

Icon	Description
	<p>Lets you determine the unit of the measurement in the graph. Select one of the following options and then click the  button for your selection to take effect:</p> <ul style="list-style-type: none"> • Default - Displays the data in its default unit, which is usually megabytes. • GB (Binary) - Displays the data in gibibytes, which is a binary unit of measurement. The computer handles the data in binary format. A gibibyte is equal to 1,073,741,824 bytes. • GB (Decimal) - Displays the data in gigabytes. A gigabyte is equal to 1,000,000,000 bytes.
	<p>Lets you format the graph to provide data within the time period specified. Select the option to the left of the Period combo box. Select one of the following options and then click the  button for your selection to take effect:</p> <ul style="list-style-type: none"> • Last Hour - Information collected in the last hour is reported. If you select Last Hour, the only frequency available is Default. • Last 24 Hours - Information collected in the last 24 hours is reported. • Last 7 Days - Information collected in the seven days is reported. • Last Month - Information collected in the last month is reported. • Last Year - Information collected in the last year is reported. • Custom - Lets you created custom settings for the graph. • Real Time - Displays the data as it is currently being gathered by the management server. The summary chart is the default chart for all real time performance display. If you select Real Time, the only frequency and unit available is Default. <p>To set a custom period, select the Custom option. See "Setting a Custom Period" on page 407 for more information.</p>

Table 111 Toolbar in Lower Pane of Performance Manager (continued)

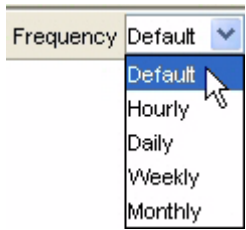






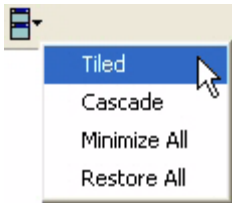
Icon	Description
	<p>Lets you change the display frequency. Select one of the following options and then click the  button for your selection to take effect:</p> <ul style="list-style-type: none"> • Default - The information is displayed according to the interval of the performance collector that gathers information for the graph. For example, if you set the collector for the graph to every 15 minutes and the period set to Last Hour, the default setting displays four data points in the graph. See "Managing Performance Collection" on page 136 for more information about setting the performance collectors. See "Creating Performance Charts" on page 400 for more information. • Hourly - The information displayed is in hourly increments. • Daily - The information displayed is in daily increments. • Weekly - The information displayed is in weekly increments. • Monthly - The information displayed is in monthly increments.
	<p>Click the  button after you have modified one or more of the following combo-boxes for your changes to take effect:</p> <ul style="list-style-type: none"> • Chart • Unit • Period • Frequency
	<p>If the chart contains more than one series of data, you can filter out the additional data series by using this Filter button.</p>
	<p>Lets you modify the performance data displayed in the graph and change graph settings.</p>
	<p>Creates a chart window.</p>

Table 111 Toolbar in Lower Pane of Performance Manager (continued)

Icon	Description
	<p>Lets you arrange the chart windows as follows:</p> <ul style="list-style-type: none">• Tiled - Displays the windows within their own tile.• Cascade - Displays the windows on top each other with the active window on top.• Minimize All - Minimizes all windows.• Restore All - Opens all windows that have been minimized.

Comparing the Performance of Different Elements

Use Performance Manager to compare the performance of different elements. Let's assume you want to compare the physical memory used on different hosts, you would do the following:

1. Access Performance Manager as described in "[Accessing Performance Manager](#)" on page 400.
2. Click the element you want to see its performance.
In this example, let's assume you clicked a host named CORTEZ.
3. Select the graph you want to see, for example, Bytes Transmitted (MB/Sec) on the Monitoring tab, located in the lower pane.
Let's assume you selected Physical Memory Used (%).
4. Scroll to the bottom of the storage tree in the lower-left pane.
5. Expand the All Elements node by clicking the (+) symbol next to the node name.
6. In the dialog window, select the elements you want to load. If you select a large number of elements, it may take some time to load these elements. When you are done, select the **Load Selected [Element Type]** button.

The software loads the information for the elements you have selected.

7. In the storage tree in the lower-left pane, expand the nodes for the hosts you have selected.
8. For each element you want to compare, select the same type of graph for each option in the left pane.

For example, if you selected "Physical Memory Used (%)" for the first element, you must select the same option for the other elements. You cannot select "Physical Memory Used (%)" for the first element and then "Processor Utilization (%)" for the second element. This is because these two options measure different types of data.

Performance Manager displays information for the different elements in the same graph.

Viewing Summary Charts


Performance Manager provides summary charts which display a single line that summarized the values for a single statistic. Multiple statistics can be shown with multiple lines.

The characteristic of the line has the following attributes:

- The vertical length of the line indicate the minimum and maximum value of the statistic within the selected data time frame.
- The green marker marks the value obtained from the latest time stamp (latest polled value).
- The purple marker marks the average value.

Summary charts are the default type for real time performance display.

To view a summary chart:

1. Access Performance Manager as described in "[Accessing Performance Manager](#)" on page 400.
2. Create a chart as described in "[Creating Performance Charts](#)" on page 400.
3. Select **Summary** from the Chart combo box in the lower pane.
4. Click the  button in the lower pane for the change to take effect.


Viewing Trending Information for Performance

The management server can display trending information in its reports. For example, you can configure Performance Manager to display trending information for the next week. This information can give you an indication of an element's future performance based on past performance.

Keep in mind the following:

- An element's performance can drastically change in the future. Please keep in mind the data trends are just assumptions and should not be treated as fact.
- Trending requires at least two sets of data gathered within the frequency specified.

To view trending information:

1. Access Performance Manager as described in "[Accessing Performance Manager](#)" on page 400.
2. Click the element for which you want to see its performance.
3. Under the Monitoring tab in the lower-left pane, select the element again. In some instances, you may need to select an element's port, such as a switch.
4. In the lower-left pane on the Monitoring tab, click a performance monitoring option.
The performance monitoring options listed in this figure vary according to the type of element. The monitoring buttons in the figure are for a switch.
5. Click the **Edit Chart Property** () button in the lower pane.
6. In the Performance Graphic Editing Dialog window, type a number in the **Trend** field.
The number corresponds to the number of frequency intervals for which the trending information will be provided. For example, if you type 5 in the **Trend** field, the chart provides trending information for five frequency intervals, such as five weeks if weeks was selected from the **Frequency** drop-down combo box menu.
7. Click **OK**.
The trending information is displayed in the chart.


NOTE: If there is not enough data to display, Performance Manager does not display the chart. For example, if you selected the weekly option from **Frequency** drop-down combo box menu and you only have two days of data, a chart is not displayed regardless of the value in the **Trend** field. Performance Manager does not display a chart if there is not enough data and the trending number is ignored.

Setting Thresholds in Performance Charts

You can set maximum and minimum thresholds in performance charts. When a value crosses a specified threshold value, the graph displays the data in red.




If a chart displays multiple statistics simultaneously, the threshold values are same for all statistics being monitored.

To set thresholds in performance charts:

1. Access Performance Manager as described in "[Accessing Performance Manager](#)" on page 400.
2. Click the element for which you want to see its performance.
3. Under the Monitoring tab in the lower-left pane, select the element again. In some instances, you may need to select an element's port, such as a switch.
4. In the lower-left pane on the Monitoring tab, click a performance monitoring option.
The performance monitoring options listed in this figure vary according to the type of element. The monitoring buttons in the figure are for a switch.
5. Click the **Edit Chart Property** () button in the lower pane.
6. In the Performance Graphic Editing Dialog window, type a value in the **Maximum** and/or **Minimum** fields. You do not need to enter a unit, such as MB, because the value entered assumes the unit of graph.
7. Click **OK**.

Removing Performance Data from a Graph

To remove multiple data from a performance graph:






1. Access Performance Manager as described in "[Accessing Performance Manager](#)" on page 400.
2. In Performance Manager, click the graph you want to modify in the bottom-right pane.
3. Click the Edit Chart Property () button in the lower pane.
4. In the Performance Graphic Editing Dialog window, click the  button corresponding with the data you want to remove from the graph.
5. When you are asked if you want to remove the data, click **Yes**.
6. Click **OK**.
7. Click the  button in the lower pane.



Setting a Custom Period

You can format the graph to provide data within a custom time period. This feature can be extremely useful for pinpointing performance changes. For example, assumed you changed the firmware of a switch two weeks ago and you want to compare the performance of that switch before and after you changed its firmware. You could create a graph that provides performance data two weeks before you changed the firmware. You could then create another graph that provides performance data two weeks after you changed the firmware.

IMPORTANT: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003/11/31, the software assumes the date is 2003/12/01.

To set a custom period:

1. Access Performance Manager as described in "[Accessing Performance Manager](#)" on page 400.
2. Click the element for which you want to see its performance.
3. Under the Monitoring tab in the lower-left pane, select the element again. In some instances, you may need to select an element's port, such as a switch.
4. In the lower-left pane on the Monitoring tab, click a performance monitoring option.
The performance monitoring options listed in this figure vary according to the type of element.
5. Select **Custom** from the Period combo box in the lower pane.
The Performance Graphic Editing Dialog window appears when you select the Custom option.
6. In the Performance Graphic Editing Dialog window, select the Custom option near the lower-left corner.
7. Click the calendar  button to the right of the **Start** field.
8. Enter the time in the time field. Make sure the time resembles a 24-hour clock, for example 22:00 for 10 p.m.
9. Click the date.
The date is highlighted in pink.
You can navigate the calendar as follows:
 -  - Displays the same month in the previous year
 -  - Displays the previous month
 -  - Displays the next month
 -  - Displays the same month in the following year
10. When you are done, click the **Set** button.
The start time and date are displayed in the **Start** field.

11. To set the end date, click the calendar  button to the right of the End field. Repeat steps 8 through 10.
- The ending time and date are displayed in the **End** field.
12. Click **OK**.
13. Click the  button.

About the Monitoring Options

See the following table for more information about some of the monitoring options.

Table 112 About the Monitoring Options

Available Monitoring Options	Available to Which Elements	Description
Active Client Logons	Microsoft Exchange Stores	The number of active client logons.
Average Delivery Time (Messages/Sec)	Microsoft Exchange Stores	The average delivery time for messages per second.
Average IO Size (Bytes/Sec)	Engenio Storage Systems	The average input/out size (Bytes/Sec)
Buffer Cache Hits Count per Second	NAS Filers (System)	The number of buffer cache hits per second.
Buffer Cache Misses Count per Second	NAS Filers (System)	The number of buffer cache misses per second.

Table 112 About the Monitoring Options (continued)

Available Monitoring Options	Available to Which Elements	Description
Buffer Hit Ratio	Oracle	The buffer cache hit ratio calculates how often a requested block has been found in the buffer cache, without requiring disk access. If the ratio is less than 90 percent investigate further. You may need to increase cache size by increasing <code>db_block_buffer(8i)</code> or <code>db_cache_size(9i)</code> .
Bytes Transferred	Engenio Storage Systems	Bytes transferred on a drive.
Bytes Transmitted (MB/Sec)	<ul style="list-style-type: none"> Storage Systems Host (port for HBA card) NAS Filer (IP Port) Switch Port 	Number of bytes transmitted by the port per second
Bytes Received (MB/Sec)	<ul style="list-style-type: none"> Storage Systems Host (port for HBA card) NAS Filer (IP Port) Switch Port 	Number of bytes received by the port per second
CPU Usage Percentage (%)	Sybase	Percentage of the Sybase CPU being used. For example, if this monitoring appears pegged at 100 percent, this means one or more Sybase databases on the host are using 100 percent of the Sybase CPU, not the host CPU. The host CPU could be pegged at 60 percent, while the Sybase CPU is pegged at 100 percent.

Table 112 About the Monitoring Options (continued)

Available Monitoring Options	Available to Which Elements	Description
CRC Errors (Errors/Sec)	<ul style="list-style-type: none"> Storage Systems Host (port for HBA card) Switch Port 	The number of cyclic redundancy checking (CRC) errors per second.
Data Received (rate/sec)	NAS Filers	The data received per second by an IP port on a filer.
Data Sent (rate/sec)	NAS Filers	The data sent per second by an IP port on a filer.
Dictionary Hit Ratio	Oracle	Use this ratio to determine shared pool is large enough to store dictionary cache data adequately. If the ratio is less than 95 percent investigate further. You may need to increase shared_pool_size.
*Disk Read (KB/second)	Disk Drives	The speed at which the disk is read. To receive this data from a 64-bit AIX host, the bos.perf.libperfstat file must be installed on the host. See the Prerequisites section in the chapter for installing the CIM Extension on AIX in the Installation Guide.
Disk Total (KB/Sec)	Disk Drives	Total speed at which the disk is read and written for HP-UX hosts. To receive this data from a 64-bit AIX host, the bos.perf.libperfstat file must be installed on the host. See the Prerequisites section in the chapter for installing the CIM Extension on AIX in the Installation Guide.

Table 112 About the Monitoring Options (continued)

Available Monitoring Options	Available to Which Elements	Description
Disk Utilization (%)	Disk Drives	The percentage of space used on the disk. To receive this data from a 64-bit AIX host, the bos.perf.libperfstat file must be installed on the host. See the Prerequisites section in the chapter for installing the CIM Extension on AIX in the Installation Guide.
*Disk Write (KB/second)	Disk Drives	The speed to which the disk is written. To receive this data from a 64-bit AIX host, the bos.perf.libperfstat file must be installed on the host. See the Prerequisites section in the chapter for installing the CIM Extension on AIX in the Installation Guide.
Exchange Services	Microsoft Exchange	Services Microsoft Exchange depends on to operate. A red line indicates the services have stopped, a green line indicates the services are running.
File Read Percent	Oracle	Percentage of "reads" for the file against the total "reads" in the database.
File Total I/O Percent	Oracle	This gives Datafile IO percentage against Total IO. Percentage of read+write for the file against total read+write in the DB.
File Write Percent	Oracle	Percentage of "write" for the file against the total "writes" in the database.
Final Destination Currently Unreachable Queue Size	Microsoft Exchange	The number of messages that cannot currently be sent to their final recipient.

Table 112 About the Monitoring Options (continued)

Available Monitoring Options	Available to Which Elements	Description
Free Physical Memory (KB)	Host	The amount of free physical memory on the host. To receive this data from a 64-bit AIX host, the bos.perf.libperfstat file must be installed on the host. See the Prerequisites section in the chapter for installing the CIM Extension on AIX in the Installation Guide.
Free Virtual Memory (KB)	Host	The amount of free virtual memory on the host. To receive this data from an AIX host, the bos.perf.libperfstat file must be installed on the host.
In Memory Sort Ratio	Oracle	This ratio shows the proportion of sorts that are performed in memory. Optimally, most sorts should be performed in memory. If the ratio is less then 90 percent investigate further. You may need to increase sort_area_size or pga_aggregate_target (for Oracle 9i if the workarea_size_policy is auto.)
Inode Cache Hits Count per Second	NAS Filers (System)	The number of inode cache hits per second.
Inode Cache Misses Count per Second	NAS Filers (System)	The number of inode cache misses per second.
Invalid CRC Errors (errors/second)	Storage Systems	The speed at which invalid cycle redundancy checking (CRC) errors are found.
Library Cache Hit Ratio	Oracle	This gives an indication of how many times parsed SQL or PL/SQL statements needed to be reloaded, which requires a reload/rebuild of the statement. If the ratio is less then 95% investigate further. You may need to increase shared_pool_size.

Table 112 About the Monitoring Options (continued)

Available Monitoring Options	Available to Which Elements	Description
Link Failures (failures/second)	<ul style="list-style-type: none">Storage SystemsHost (port for HBA card)Switch Port	The number of link failures per second.
Mail Box Count	Microsoft Exchange Stores	The number of mail boxes for a store.
Memory Usage Percentage (%)	Sybase	Percentage of the Sybase memory being used, not the host memory.
Messages Awaiting Directory Lookup Queue Size	Microsoft Exchange	The number of messages waiting for the recipient to be resolved in the global catalog, so the message can be sent.
Messages to be Routed Queue Size	Microsoft Exchange	The number of messages waiting to be routed to the recipient.
Name Cache Hits per Second	NAS Filers	The number of cache hits per second.
Name Cache Misses per Second	NAS Filers	The number of name cached misses per second on a NAS filer.
Packets Received (rate/second)	NAS Filers (IP Port)	The number of packets received per second by the IP ports on a filer.

Table 112 About the Monitoring Options (continued)

Available Monitoring Options	Available to Which Elements	Description
Packets Sent (rate/sec)	NAS Filers (IP Port)	The packets sent per second by the IP ports on a filer.
Parse CPU to Total CPU Ratio	Oracle	Ratio closer to 0 percent is good. A high ratio means system is performing too many parses and indicates that shared pool is configured poorly or application is using SQL and so SQL cannot be shared. Setting Cursor_sharing=Force or SIMILAR might help in this case
Pre-submission Queue Size	Microsoft Exchange	The number of messages being held by the Pre-submission Queue, which holds messages waiting to be submitted to the Messages to be Routed Queue.
Processor Utilization (%)	Hosts	The percentage of the processor being used. To receive this data from a 64-bit AIX host, the bos.perf.libperfstat file must be installed on the host. See the Prerequisites section in the chapter for installing the CIM Extension on AIX in the Installation Guide.
Physical I/O Percentage (%)	Sybase	Percentage of the Sybase physical/input output, not the host physical/input output.
Physical Memory Used (%)	Hosts	The percentage of physical memory used on the host. To receive this data from a 64-bit AIX host, the bos.perf.libperfstat file must be installed on the host. See the Prerequisites section in the chapter for installing the CIM Extension on AIX in the Installation Guide.
Percent Read (%)	Engenio Storage Systems	The percentage read.

Table 112 About the Monitoring Options (continued)

Available Monitoring Options	Available to Which Elements	Description
Processor Busy Time	NAS Filers (processor)	The amount of time the processor is busy.
Processor Elapsed Time	NAS Filers (processor)	The amount of time that has passed since the NAS filer was rebooted.
Read IO Rate (Reads/Sec)	Engenio Storage Systems	The input/output of the read rate.
Read Operations	EMC and Engenio Storage Systems	Read operations (bytes/second).
Read Requests	EMC and Engenio Storage Systems	Read requests (bytes/second).
Receive Queue Size	Microsoft Exchange Stores	The size of the receive queue.
Recovered Errors	Engenio Storage Systems	The number of recovered errors on the drive.
Redo Buffer Allocation Retries Ratio	Oracle	It is number of times DSP had to wait to write to log buffer. If the ratio is more then 1 percent investigate further and consider increasing redo log buffer size.
Redo Logspace Request Ratio	Oracle	The number of times lgwr had to wait for writing to redo log file. If the ratio is more then .0002 (1 in 5000) investigate further and consider increasing redo log buffer size.
Requests Serviced	EMC and Engenio Storage Systems	Requests serviced (bytes/second).
Reserved Inodes	NAS Filer (File System)	The number of reserved inode files on a NAS file system.

Table 112 About the Monitoring Options (continued)

Available Monitoring Options	Available to Which Elements	Description
Retried Requests	Engenio Storage Systems	The number of retried requests for a drive.
Send Queue Size	Microsoft Exchange Stores	The size of the send queue.
SMTP Local Delivery Queue Size	Microsoft Exchange	The number of messages in the SMTP Local Delivery Queue, the queue hold mail box of local recipients.
SMTP Server Queues Summary	Microsoft Exchange	Provides a count of the total number of messages in the following queues: <ul style="list-style-type: none"> • Final Destination Currently Unreachable • Messages Awaiting Directory Lookup • Messages to Be Routed • Pre-submission • SMTP Local Delivery
Storage Group Size	Microsoft Exchange Storage Groups	The size of the storage group in megabytes.
Store size (MB)	Microsoft Exchange Stores	The size of the store in megabytes.
System Event Time Waited (ms)	Oracle	Displays the delta of time waited for the system wait events.
Tablespace Read Percent	Oracle	Percentage of “reads” for the tablespace against the total “reads” in the database.
Tablespace Write Percent	Oracle	Percentage of “writes” for the tablespace against the total “writes” in the database.
Tablespace Total I/O Percent	Oracle	This gives the tablespace input/output percentage against the total input/output. Percentage of read+write for the tablespace against total read+write in the database.

Table 112 About the Monitoring Options (continued)

Available Monitoring Options	Available to Which Elements	Description
Timeouts	Engenio Storage Systems	The number of timeouts on the drive.
Total Bandwidth (Bytes/Sec)	Engenio Storage Systems	The Total bandwidth.
Total Inodes	NAS Filers (File systems)	The total number of inode files for a NAS file system.
Total IOs	Engenio Storage Systems	The total input/output.
Total IO Rate (IOs/Sec)	Engenio Storage Systems	The total input/output rate.
Unrecovered Errors	Engenio Storage Systems	The number of unrecovered errors on a drive.
Used Inodes	NAS Filers (File Systems)	The number of unused inode files on a NAS file system.
Virtual Memory Used (%)	Hosts	The percentage of virtual memory used on the host. To receive this data from an AIX host, the bos.perf.libperfstat file must be installed on the host.
Write IO Rate (Writes/Sec)	Engenio Storage Systems	The input/output of the write rate.
Write Operations	EMC and Engenio Storage Systems	Write operations (bytes/second).

*These monitoring options are not available on HP-UX hosts because HP-UX hosts do not return read/write data separately.

Task Already Running on Element

If the software is trying to collect information for monitoring or statistic collection and the element is not returning the information in time, you are shown a message resembling the following:

```
Task xxxx is already running on element xxxx
```

The software cannot obtain the information in a timely manner because one or more of the following:

- The element might be inherently slow.

- The element might be busy with other tasks.
- You are trying to collect too much information at once from the element.

If you think you might be trying to collect too much information from the element, you might want to narrow down the collection. For example, if you are trying to collect monitoring information for three disks drives on a server, you might want to try collecting information for one disk drive.

The software stops trying to gather information after several unsuccessful attempts.

Monitoring with Direct Attached Storage

A port on a storage system directly attached to a host, does not appear in the left panel for monitoring. If you want to monitor the port, connect the port to a switch.

Supported Host Configurations for Monitoring

Refer to the following table to determine which host configurations the management server can monitor. The management server supports configurations that have a Y next to them. Nonsupported configurations have an N next to them.

NOTE: If a configuration listed in the following table is not supported, you can still obtain processor and memory statistics from the host. The exception, however, is Windows NT 4, which does not provide any monitoring information to the management server.

If the host has several configurations listed in the table below and one of them cannot be monitored, monitoring is not supported for all of the configurations on the host. For example, assume you have a host with Solaris 9 Sun Foundation Suite Leadville with MPXIO and Solstice DiskSuite/Volume Manager. Even though the management server supports monitoring for Solstice DiskSuite/Volume Manager, both of those devices, Sun Foundation Suite Leadville with MPXIO and Solstice DiskSuite/Volume Manager, cannot be monitored because Solaris 9 Sun Foundation Suite Leadville with MPXIO is not supported, as shown with the following formula:

$$\begin{array}{l} \text{A monitorable configuration (Y) + a unmonitorable configuration (N) =} \\ \text{unmonitorable configuration (N)} \end{array}$$

Keep in mind the following:

- In all configurations, you cannot monitor a VCM database device.
- The management server only monitors the top or bottom layer of Solstice DiskSuite/Volume Manager. For example assume you have a normal configuration for Solstice DiskSuite/Volume Manager (/folder <- d1 <- d2 <- d3 <- d4 <- cxydzs#). The management server reports on the folder and the highest layer, which is d1. Assume you have soft partitioning (/folder <- d1 <- d2 <- d3 <- d4 <- cxydzs#). The management server reports on the folder and the lowest layer, which is d4. The management server skips all layers between the d# devices. You are not able to monitor the middle layers, which are d2 and d3 in the previous example.

The following table is a sample of possible configurations, not a complete list of all products in the marketplace.

Table 113 Host Monitoring Support

Host Configuration	Monitoring Supported?
Solaris 8,9	Y
Solaris 8,9 VXVM	Y
Solaris 8,9 VXVM dmp	Y
Solaris 8,9 PowerPath	Y
Solaris 8,9 HDLM	N
Solaris 8,9 RDAC	N
Solaris 8,9 Sun Foundation Suite Leadville	Y
Solaris 8,9 DAS	N
Solaris 8,9 Sun Foundation Suite Leadville + MPXIO	N
Solaris 8,9 Solstice DiskSuite/Volume Manager	Y
Solaris 8,9 Sun Foundation Suite Leadville + MPXIO + Solstice DiskSuite/Volume Manager	N
Windows NT 4	N
Windows 2000	Y
Windows 2000 PowerPath	N
Windows 2000 HDLM	N
Windows 2000 RDAC	N
Windows 2000 DAS	N
Windows 2000 Volume Manager	N
Windows 2000 Volume Manager and PowerPath	N
Windows 2000 Volume Manager and HDLM	N
SGI ProPack 3.0	Y
SGI ProPack 3.0 XVM	Y
SGI ProPack 3.0 CXFS	Y (only on node sending I/O)
AIX 5.1, 5.2 with LVM	Y
AIX 5.1 with PowerPath	Y
AIX 5.1 with HDLM	Y

Table 113 Host Monitoring Support (continued)

Host Configuration	Monitoring Supported?
AIX 5.3	Y
Irix 6.5.x	Y
Irix 6.5.x XVM	Y
Irix 6.5.x CXFS	Y (only on node sending I/O)
Redhat 2.1	Y
Redhat 3.0 with Sistina LVM	Y
Redhat 2.1 with HDLM	N
Redhat 3.0	Y
Redhat 3.0 with HDLM	N
HP-UX 11i	Y
HP-UX 11.0	Y
HP-UX 11i LVM	Y
HP-UX 11.0 LVM	Y
HP-UX 11i PV Link	Y
HP-UX 11.0 PV Link	Y
HP-UX 11i PV Link Volumes	Y
HP-UX 11.0 PV Link Volumes	Y
HP-UX 11.i with PowerPath	Y
HP-UX 11.0 with PowerPath	Y
HP-UX 11.0 with HDLM	Y
HP-UX 11.i with HDLM	Y

Sudden Dips Displayed in Certain Charts in Performance Manager

In Performance Manager and on the Monitoring tab, charts that display data gathered by certain volume and drive counters display their charts with the results increasing to the maximum value and then decreasing rapidly to a very low number that starts rising again.

The sudden dip in the charts occur because the counters holding the values displayed in the charts can contain at most 231 (2147483648) integers. The counters revert back to zero after reaching 231 (2147483648) integers. Then, they continue to go up again. The counters do not usually display zero in a chart because they are quickly gathering data again.

The following charts for individual drives are impacted: ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, Unrecovered Errors, Recovered Errors, Timeouts, Retried Requests

The following charts for individual volumes are impacted: Bytes Read, Bytes Read Large, Bytes Written, Bytes Written Large, ReadIOs Large, WriteIOs Large, ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, ReadHitIOs

Values Continue to Increase in Charts for Aggregated Drives and Aggregate Volumes

Values in charts for aggregate drives and aggregate volumes continue to rise smoothly in Performance Manager and on the Monitoring tab. The only time the values in the charts come back down to the values returned by the array, is when the service for the management server restarts. Then, the charts display the values returned by the array. The values in the charts continue to rise until the service for the management server is stopped.

The values in the charts continue to rise because the values do not revert back to zero when an individual drive or volume counter that makes up the aggregation contain 231 (2147483648) integers. Instead, the values continue to rise smoothly.

The following charts for aggregate drives are impacted and charts for aggregated drives are also provided for arrays, controllers, and volume groups: ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, Unrecovered Errors, Recovered Errors, Timeouts, Retried Requests

The following charts for aggregate volumes are impacted, and charts for aggregated volumes are also provided for arrays, controllers, and volume groups: Bytes Read, Bytes Read Large, Bytes Written, Bytes Written Large, ReadIOs Large, WriteIOs Large, ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, ReadHitIOs

14 Finding an Element's Storage Capacity

IMPORTANT: Depending on your license, Capacity Manager may not be available. See the "List of Features" to determine if you have access to Capacity Manager. The "List of Features" is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).

This chapter describes the following:

- ["About Capacity Manager"](#) on page 423
- ["Accessing Capacity Manager"](#) on page 425
- ["About the Toolbars in Capacity Manager"](#) on page 425
- ["Finding the Capacity of an Element"](#) on page 428
- ["Obtaining Utilization Reports"](#) on page 430
- ["Printing Elements in Capacity Manager"](#) on page 431
- ["Viewing Capacity Charts"](#) on page 433
- ["Viewing Trending Information for Storage Capacity"](#) on page 435
- ["Different Results for the df -k Command and Capacity Manager"](#) on page 435

About Capacity Manager

IMPORTANT: Depending on your license, Capacity Manager may not be available. See the "List of Features" to determine if you have access to Capacity Manager. The "List of Features" is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).

Capacity Manager provides graphical representation of an element's storage capacity and utilization in the storage network. It provides the following tabs:

- **List**
- **Path**
- **Utilization**
- **Capacity Data**
- **Capacity Chart**

Keep in mind the following:

- The Capacity Manager displays the total capacity of an application, including the network drives. If you look at the capacity of the application in Chargeback Manager, the capacity differs. Chargeback Manager provides only network capacity. See ["Capacity in Chargeback Manager and Capacity Manager Differ"](#) on page 476 for more information.
- Volume names from ambiguous automounts on Solaris hosts are not displayed in Capacity Manager. See ["Volume Names from Ambiguous Automounts Are Not Displayed"](#) on page 508.

About the List Tab

The List tab lets you quickly access an element. For example, to quickly access a host in the topology, expand the List View Tree node, then select your host in the tree by expanding the All Elements and Hosts nodes. When you select an element in the tree, it appears highlighted in the topology.

About the Path Tab

The Path tab provides information about an element's path. By clicking a host's node, you can determine the host's path in the application. When you click a host node in the tree, the elements in the host's path appear highlighted in the right pane.

About the Utilization Tab

Provides host and switch reports about certain aspects of utilization and storage, for example:


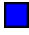

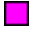
- **Remote vs. Local Utilization**
- **Remote Storage**
- **Local Storage**
- **Total Port Utilization**
- **Port Utilization by Fabric**

See ["Obtaining Utilization Reports"](#) on page 430 for more information.

About the Capacity Data Tab

The Capacity Data Tab provides information about the capacity of an element. You can quickly determine the general capacity of an element by its colors, as explain in the following legend.

Figure 72 Legend for Capacity Manager

Color	Description
	Description varies according to element type: <ul style="list-style-type: none">• For storage systems - The space is mapped.• For all elements except storage systems - The space is used.
	Unallocated
	Unused Raw (storage systems only)
	Unmapped (storage systems only)

For example, the element displayed in the following figure is about 75 percent available. The rest if it is being used.



Figure 73 Capacity of an Element

You can obtain more detailed information about an element by clicking it in the right pane or in the Capacity Manager tree, as explained in the topic, See ["Finding the Capacity of an Element"](#) on page 428 for more information.

About the Capacity Chart Tab

The Capacity Chart tab lets create bar or line charts to view your capacity data. You can use these charts to display trending information. See ["Viewing Capacity Charts"](#) on page 433 and ["Viewing Trending Information for Storage Capacity"](#) on page 435 for more information.

Accessing Capacity Manager

To access Capacity Manager, click **Optimize > Storage Essentials > Capacity Manager**.

About the Toolbars in Capacity Manager

Capacity Manager provides two toolbars. One in the upper pane and another in the lower pane. The toolbar in the upper pane is the same as the one in System Manager. See ["About the Toolbar in System Manager"](#) on page 164 for information about the toolbar in the upper pane.

When the Capacity Chart tab is active, the toolbar in the lower pane provides the following information:

Table 114 Toolbar in Lower Pane of Capacity Manager


Icon	Description
	<p>Lets you switch between the following:</p> <ul style="list-style-type: none">• Bar - Displays each data point as a bar. The data for the different elements is displayed side by side.• Line - Displays each data point as a dot with a line connected to the previous data points. The data for the different elements for a specific point in time is displayed in the same column.

Table 114 Toolbar in Lower Pane of Capacity Manager (continued)

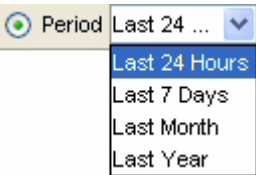

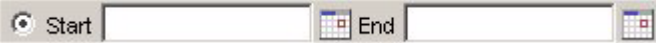


Icon	Description
	<p>Lets you format the graph to provide data within the time period specified. Select the option to the left of the Period combo box. Select one of the following from the drop-down combo box. Then, click the  button to update the chart.</p> <ul style="list-style-type: none">• Last 24 Hours - Information collected in the last 24 hours is reported. This option is only available to hosts and applications.• Last 7 Days - Information collected in the last seven days is reported.• Last Month - Information collected in the last month is reported.• Last Year - Information collected in the last year is reported.
	<p>Lets you format the graph to provide data within the starting and ending time specified. Select the option to the left of the Start field. Click the  icon. In the Time field, type the time you want the graph to start. Use the 24-hour clock. Select a date and click Set. Repeat the steps for the End field. Then, click the  button to update the chart.</p> <p>Important: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the first day of the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.</p>

Table 114 Toolbar in Lower Pane of Capacity Manager (continued)

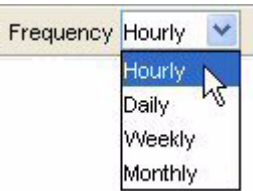
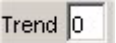




Icon	Description
	<p>(Applications and Hosts only) Lets you change the display frequency. The options are the following:</p> <ul style="list-style-type: none"> • Hourly - The information is displayed in hourly increments. • Daily - The information is displayed in daily increments. • Weekly - The information is displayed in weekly increments. • Monthly - The information is displayed in monthly increments.
	<p>Lets you set trending information. In the pane displaying the chart, type an integer in the Trend field. The integer corresponds to the number of frequency intervals for which the trending information will be provided. For example, if you type 5 in the Trend field, the chart provides trending information for five frequency intervals, such as five weeks if weeks was selected from the Frequency drop-down combo box menu. Click the  button to update the chart.</p> <p>When a switch or storage system is selected, the frequency field is set to hourly.</p> <p>See "Viewing Trending Information for Storage Capacity" on page 435.</p> <p>Keep in mind the following:</p> <ul style="list-style-type: none"> • An element's performance can drastically change in the future. Please keep in mind the data trends are just assumptions and should not be treated as fact. • Trending requires at least two sets of data gathered within the frequency specified.

Table 114 Toolbar in Lower Pane of Capacity Manager (continued)

Icon	Description
	Applies period, frequency and trending information.
	If the chart contains more than one series of data, you can filter out the additional data series by using this Filter button.
	Lets you print a graph.

Finding the Capacity of an Element

NOTE: Capacity Manager rounds the data it displays. This rounding could result in discrepancies in the totals you add for a property against the data displayed in the Summary column. For example, if you add the total capacity from each data pool and compare that total against the number for Total Capacity displayed under the Summary heading, you will most likely find the results differ.

To find the capacity of an element:

1. Access Capacity Manager as described in “[Accessing Capacity Manager](#)” on page 425.
2. (Optional) To quickly view the capacity of all of the elements in a fabric or application path, click the fabric or application path displayed in the tree for Capacity Manager, as shown in the following figure.




Figure 74 Viewing the Capacity of Elements in a Fabric

- See “[About Capacity Manager](#)” on page 423 for an explanation of the colors displayed.
3. Do one of the following:
 - Click an element in the right pane.
 - Click an element in the tree for Capacity Manager.

When an element is clicked, a pane appears along the bottom of the page and it displays information under the **Capacity Summary** and **Capacity** columns (**Port Summary** and **Port** columns for switches), such as the following for applications, hosts and storage systems:

- **Total Capacity** - If a switch is selected, the total number of ports appear instead.
- **Total Used**
- **Available**
- ***Unmounted Volume**
- **Percentage used**

*Appears only for hosts and filers. The Unmounted Volume field under Capacity Summary automatically displays 0 MB if you discovered the host but not the storage system connected to it. This may occur if you did not enter the IP address of the storage system when performing discovery and/or your license does not allow you to discover a particular storage system. See the Supported Elements section in the “List of Features” to determine which storage systems you can discover. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).

The following is displayed for storage systems under the **Capacity** and **Capacity Summary** columns:

- **Total Capacity**
- **Mapped**
- **Unmapped**
- **Total Allocated**
- **Unallocated**
- **Unused Raw***
- **Percentage Allocated***

*You may need to use the scroll bar in the pane to view these properties.

Additional Information for Applications

The following additional information is displayed for storage groups (Microsoft Exchange) or database files (Oracle):

- **Total Capacity**
- **Total Used**
- **Available**
- **Percentage Used** - The percentage used compared to the total capacity of the storage group or database file.

Additional Information for Hosts and Filers

The following additional information is displayed for each storage volume:

- **Total Capacity**
- **Total Used**
- **Available**

- **Percentage Used** -The percentage used compared to the total capacity of the storage volume.

Additional Information for Storage Systems

Capacity Manager displays additional information for each storage pool (Volume Group):

- **Total Capacity** - Total allocated, unallocated and unused raw.
- **Mapped** - Provisioned volumes exported to a system port. Mapped volumes are accessible.
- **Unmapped** - Provisioned volumes that are not exported to a system port. Unmapped volumes are not accessible.
- **Total allocated** - Provisioned volume.
- **Unallocated** - Available space left on the storage pool
- **Percentage Allocated** - The percentage allocated compared to the total capacity of the storage system.

NOTE: To sort information by ascending order in a column, click the column heading.

The following table describes how the properties are calculated for HDS array groups.

Table 115 Explanation of the Properties of the Capacity Levels for HDS Array Groups

Property Displayed for an Array Group (CIM_StoragePool)	How It Is Calculated	Explanation
Total Capacity	Sum of volume.size + pool.totalRemainingSpace	Sum of sizes of all LDEVs in the array group plus the total free space
Mapped	Sum of volume.size for each volume that has a LUN	LDEVs that have LUNs
Unmapped	Sum of volume.size for each volume that doesn't have a LUN	LDEVs that don't have LUNs but are not on the management server's "free" list
Total Allocated	Sum of volume.size for all volumes	All LDEVs that are not on the management server's "free" list
Unallocated	pool.totalRemainingSpace	LDEVs on the management server's "free" list plus the total free space

Obtaining Utilization Reports

The software provides the following utilization reports to help you determine how much of your storage is being used:

- **Host Storage Utilization Reports**
- **Switch Utilization Reports**
- **Subsystem Utilization Reports**

NOTE: Capacity Manager does not provide utilization reports for NetApp filers.

To view a utilization report:

1. Access Capacity Manager as described in "[Accessing Capacity Manager](#)" on page 425.
2. Click the **Utilization** tab.
3. To view a utilization report, click one of the reports listed in the tree, as shown in the following figure.






Figure 75 Viewing a Utilization Report

The report appears in the right pane.

Printing Elements in Capacity Manager

The software lets you print the topology in Capacity Manager. For example, you can provide a printout to upper management that shows not only the topology of the network, but also the capacity of each element.

To print the elements in Capacity Manager:

1. Access Capacity Manager as described in "[Accessing Capacity Manager](#)" on page 425.
2. If the topology spans more than one screen, arrange the elements so they are closer together, preferably on one screen. To move an element, click the  button and then the element you want to move. Drag the element to its new location. Moving elements closer together prevents the printout from appearing too stretched.
3. Click the  button.
4. Use the fields on the Paper tab to modify the setup of the page. When you are done, click **Apply**. If you want the default settings, click **Default**.
A preview of the printout is displayed in the right pane.

IMPORTANT: Before you change the margins, decide on a unit of measurement.

- **Paper format** - Select the paper size from the drop-down menu.
 - **Unit** - Select cm (centimeters) or inch for the margins.
 - **Paper width** - Displays the width of the paper. You can modify the measurement in this field when you select the **Custom** option in the Paper format drop-down menu.
 - **Paper height** - Displays the height of the paper. You can modify the measurement in this field when you select the **Custom** option in the Paper format drop-down menu.
 - **Top margin** - Type a measurement for the top margin.
 - **Bottom margin** - Type a measurement for the bottom margin.
 - **Left margin** - Type a measurement for the left margin.
 - **Right margin** - Type a measurement for the right margin.
 - **Orientation** - Click an orientation for the printout.
5. Click the **View Selection** tab to modify how the printout will appear on the page. You can modify the following. When you are done, click **Apply**. If you want the default settings, click **Default**.
- A preview of the printout is displayed in the right pane.

IMPORTANT: Before you change the margins, decide on a unit of measurement.

- **Start x** - Determines the horizontal placement of the printout on the page with zero being the closest to the right margin. For example, if the value is 50 for **Start x**, the printing starts at 50 inches or centimeters (depending on what you selected) from the right margin. You can also enter negative numbers. Anything more than zero expands the printout to another page.
 - **Start y** - Determines the vertical placement of the printout on the page with zero being the closest to the bottom margin. For example, if the value is 50 for **Start y**, the printing starts at 50 inches or centimeters (depending on what you selected) from the bottom. You can also enter negative numbers.
 - **Width** - Determines the width of the printout.
 - **Height** - Determines the height of the printout.
- To remove extra space around the topology, click the **Trimmed** button.
6. To change how many pages the printout will use, select one of the following. When you are done, click **Apply**. If you want the default settings, click **Default**.
- A preview of the printout is displayed in the right pane.

IMPORTANT: Before you change the margins, decide on a unit of measurement.

- **Unit** - Select cm (centimeters) or inch for the margins.

- **Position/Size** - Lets you change the position and size of the printout so that it spans several pages:
 - **Start x** - Determines the horizontal placement of the printout. Anything more than zero expands the printout to another page.
 - **Start y** - Determines the vertical placement of the printout.
 - **Width** - Determines the width of the printout. If the width entered does not fit on the page, the printout wraps around to another page.
 - **Height** - Determines the height of the printout. If the height entered does not fit on the page, the printout wraps around to another page.
 - **Resolution (pixel/unit)** - Lets you change the resolution so that the printout spans several pages.
 - **Page** - Lets you expand the printout so it prints on several pages without modifying the chart.
7. To preview your pages, click the Preview tab. Then click the page you want to preview. The page appears in the right pane.
 8. When you are ready to print, click **Print**.
 9. Click **Close**.

NOTE: To revert back to all of the original settings, click the **Default** button next to the **Print** button.

Viewing Capacity Charts

Capacity Manager provides a graphical representation of the capacity history of an element, such as port summary information for switches.

You can manipulate the charts, so they show a different reporting period and frequency. For example, you could show the capacity of a host over the past 24 hours with an hourly monitoring frequency.

Keep in mind the following:





- Verify the performance collector for the element is enabled (**Optimize > Storage Essentials > Performance Data Collection**). See "[Managing Performance Collection](#)" on page 136 for more information about enabling performance collectors.
- Switches and storage systems display data from the last time Discovery Data Collection was performed. See "[Updating the Database with Element Changes](#)" on page 45 for more information.
- If you see "There is not enough data to produce a chart [chart_title] at this time." message, lessen the frequency option, so that the amount of time listed in the **Frequency** drop-down menu has passed before you view the chart again. For example, if the **Frequency** menu displays hourly, you need to wait an hour for data to appear in the chart.


To find the performance of an element:

1. Access Capacity Manager as described in "[Accessing Capacity Manager](#)" on page 425.



2. Click the element for which you want to see a history of its capacity.
3. In the bottom split pane, click the **Capacity Chart** tab.
4. In the lower-middle split pane, click a report title.
The chart for the monitoring option is displayed in the lower-right pane.

NOTE: You can change the orientation of the chart by holding down the mouse button when you click the chart and moving the mouse.

5. To change the reporting period, do one of the following. Then, click the  button to update the chart.
 - **Display data within a time period** - You can format the graph to provide data within the time period specified. Select the option to the left of the **Period** combo box. Select one of the following from the drop-down combo box. Then, click the  button to update the chart.
 - **Last 24 hours** - Information collected in the last 24 hours is reported. This option is not available to switches and storage systems.
 - **Last 7 Days** - Information collected in the seven days is reported.
 - **Last Month** - Information collected in the last month is reported.
 - **Last Year** - Information collected in the last year is reported.
 - **Display data within a starting and ending time** - You can format the graph to provide data within the starting and ending time specified. Select the option to the left of the **Start** field. Click the  icon. In the **Time** field, type the time you want the graph to start. Use the 24-hour clock. Select a date and click **Set**. Repeat the steps for the **End** field. Then, click the  button to update the chart.

Important: If you change the date in the field to a date that does not exist in a month, the software automatically calculates the date to the first day of the next month. For example, if you enter 2003-11-31, the software assumes the date is 2003-12-01.
6. (Applications and Hosts only) To change the display frequency, select one of the following from the **Frequency** drop-down combo box.
 - **Hourly** - The information is displayed in hourly increments.
 - **Daily** - The information is displayed in daily increments.
 - **Weekly** - The information is displayed in weekly increments.
 - **Monthly** - The information is displayed in monthly increments.
7. If the chart contains more than one series of data, you can filter out the additional data series by clicking the  button.
8. To add trending information, type an integer in the **Trend** field.
The integer corresponds to the number of frequency intervals for which the trending information will be provided. For example, if you type 5 in the **Trend** field, the chart provides trending information for five frequency intervals, such as five weeks if weeks was selected from the **Frequency** drop-down combo box.

See “[Viewing Trending Information for Storage Capacity](#)” on page 435.

9. Click the  button to update the chart.
10. To print the chart, click the  button displayed in the same pane as the chart.

Viewing Trending Information for Storage Capacity

The management server can display trending information in its reports. For example, you can configure Capacity Manager to display trending information for the next week. This information can give you an indication of an element's future capacity need based on its past capacity utilization.

Keep in mind the following:

- An element's capacity can drastically change in the future. Please keep in mind the data trends are just assumptions and should not be treated as fact.
- Trending requires at least two sets of data gathered within the frequency specified.

To view trending information:

1. Access Capacity Manager as described in “[Accessing Capacity Manager](#)” on page 425.
2. Click the element for which you want to see a history of its performance.
3. In the bottom pane, click the **Capacity Chart** tab.
4. In the lower-middle split pane, click a monitoring option.
5. In the pane displaying the chart, type a number in the **Trend** field.

The number corresponds to the number of frequency intervals for which the trending information will be provided. For example, if you type 5 in the **Trend** field, the chart provides trending information for five frequency intervals, such as five weeks if weeks was selected from the **Frequency** drop-down menu.

6. Click the  button.

The trending information is displayed.

NOTE: If there is not enough data to display, Capacity Manager does not display the chart. For example, if you selected the weekly option from **Frequency** drop-down menu and you only have two days of data, a chart is not displayed regardless of the value in the **Trend** field. Capacity Manager does not display a chart if there is not enough data and the trending number is ignored.

Different Results for the `df -k` Command and Capacity Manager

If you run the `df -k` command on UNIX for a storage system, you may notice the total capacity displayed does not match the total capacity in Capacity Manager. This difference occurs because Capacity Manager calculates the total capacity differently than the `df -k` command. The `df -k` command calculates the total capacity as follows:

```
used capacity + available capacity + reserved capacity
```

Capacity Manager calculates the capacity differently, as shown below:

`used capacity + available capacity`

The difference between the two calculations is the capacity reserved for superuser. If a file system has a reserved capacity, the total capacity from the `df -k` command and Capacity Manager will differ.

For example, assume you ran the `df -k` command for the file system, `/dev/dsk/c0t0d0s0`. After you run the `df -k` command, you notice that the total capacity displayed is 6688076 KB. When you look at Capacity Manager, the total capacity displayed is 6621196 KB. Actually, Capacity Manager displays results in megabytes, but for this example, it is easier to have the totals using the same units.

The totals differ. How did this happen? When you ran the `df -k` command, the computer ran the equation mentioned earlier (`used capacity + available capacity + reserved capacity`):

`1904031 KB + 4717165 KB + 66880 KB = 6688076 KB`

where

- 1904031 KB is the used capacity
- 4717165 KB is the available capacity
- 66880 KB is the capacity reserved for the superuser. The percentage of the reserved capacity can be set when using the `newfs -m` command.

Capacity Manager calculated the total capacity by using the equation discussed previously (`used capacity + available capacity`) and displayed the result in megabytes:

`1904031 KB + 4717165 KB = 6621196 KB`

where

- 1904031 KB is the used capacity
- 4717165 KB is the available capacity

Because Capacity Manager does not include the reserved capacity in its calculations, the difference between the two calculations is the capacity reserved for the superuser, which is 66880 KB.

15 Managing Policies

IMPORTANT: Depending on your license, Policy Manager may not be available. See the “List of Features” to determine if you have access to Policy Manager. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).

This chapter describes the following:

- “[About Policy Manager](#)” on page 437
- “[Accessing Policy Manager](#)” on page 438
- “[Creating Policies](#)” on page 438
- “[Modifying Policies](#)” on page 446
- “[Viewing Policies](#)” on page 453
- “[Deactivating a Policy](#)” on page 453
- “[Deleting Policies](#)” on page 453
- “[Providing E-mail Notification for a Policy](#)” on page 454
- “[Providing Event Generation for a Policy](#)” on page 454
- “[Providing a Custom Command for a Policy](#)” on page 455

About Policy Manager

Policy Manager can automatically send an e-mail, generate an event, or run a custom script when an element is being overused or when one of the following occurs:

- A new element is discovered
- Successful provisioning occurred
- An event occurred on one or more specified elements

IMPORTANT: Policies are not triggered when a collector is running. For example, assume you configured a policy to be triggered when a host has a critical event. If a collector is running when a critical event occurs, the policy is not triggered. Collectors are used to gather information for reports, monitoring and File System Viewer. .

if the collector for a report or Performance Manager is running, the policy is not triggered.

Policy Manager lets you create the following types of policies:

- **Utilization policies** - Monitors the utilization of an element. Options provided depend on the type of element.
For example, you can configure Policy Manager so you receive an e-mail message when the amount of free space on a server decreases to a specified level.
- **Infrastructure policies** - Monitors the following:

- A new element is discovered
- Successful provisioning occurred
- An event occurred on one or more specified elements

Accessing Policy Manager

This topic provides the various techniques for accessing Policy Manager. You have only one way to access policies for discovery, provisioning, and events; however, you have multiple ways to access policies. The flexibility in accessing utilization policies lets you easily create and manage policies without interrupting your work flow.

Policy Manager provides the following options for accessing policies:

- **To access from System Manager** - Double-click an element in System Manager and then click the **Policies** tab. This method displays the utilization policies for just the element that was double-clicked.
- **To access policies for file servers** - Access the Policy tab in Application Viewer. Expand the **Applications** and **File Servers** nodes in the tree in Application Viewer. Click the file server name in the Application Viewer tree. Click the **Policies** tab in the right pane.

To access all policies in Policy Manager:

1. Click **Tools > Storage Essentials > Policy Manager**.
2. Click a policy in the Policy Manager tree.

Creating Policies

This section describes the following:

- "[Creating a Utilization Policy](#)" on page 438
- "[Creating Policies for Discovery](#)" on page 441
- "[Creating Policies for Provisioning](#)" on page 442
- "[Creating Policies for Events](#)" on page 444
- "[Testing a Utilization Policy](#)" on page 446

Creating a Utilization Policy

You can create a utilization policy that generates an event, sends an e-mail, or runs a custom command when an element is being overused. For example, you can configure Policy Manager so you receive an e-mail message when the amount of free space on a server decreases to a specified level.

Keep in mind the following:

- If you plan to use e-mail notification with your policy, first assign an SMTP server from which the management server can send its e-mail notifications. See "[Setting Up E-mail Notification](#)" on page 99.
- Policies that are triggered for virtual applications are also triggered for file servers.

To create a utilization policy:

1. Access Policy Manager as described in the topic, "[Accessing Policy Manager](#)" on page 438.
2. In the left pane, select an element or element type you want the policy to apply to.
3. In the right pane, click the **Add** button.
4. Select a policy. Then, click **Next**.

The following types may be displayed.

Table 116 Policy Templates

Policy Type	Description	Applies to
Free Space	Amount of free space on one of the following: <ul style="list-style-type: none"> • A host • A Microsoft Exchange instance • A database instance, such as Microsoft SQL Server, Sybase, or Oracle. • File Servers 	<ul style="list-style-type: none"> • Hosts • Backup clients • Microsoft Exchange • Microsoft SQL Server • Oracle • Sybase • File Servers
Percent Used	Description varies according to element type: <ul style="list-style-type: none"> • Hosts - Percent of storage used for a host. • Microsoft Exchange instances - Percent of storage used for a Microsoft Exchange instance. • Database instances - Percent of storage used for a database instance, such as Microsoft SQL Server, Sybase, or Oracle. • Switches - Percent of ports used for a switch. • File Servers - Percentage of disk space used. 	<ul style="list-style-type: none"> • Hosts • Backup clients • Microsoft Exchange • Microsoft SQL Server • Oracle • Sybase • Storage Systems • Switches • File Servers
Unmapped Storage	Space not assigned to any mapped volume for the storage system	Storage systems
Unmapped Storage Percent	Percent of total space not assigned to any mapped volume for the storage system	Storage systems
Unused Raw Storage	Space not assigned to any storage pool for a storage system.	Storage systems
Unused Raw Storage Percent	Percent of total space not assigned to any storage pool for a storage system.	Storage systems
Unused Ports	Number of unused ports for a switch or storage system.	<ul style="list-style-type: none"> • Switches • Storage systems
Port Usage	Percent of ports used for a storage system.	Storage Systems

5. In the **Name** field on the Policy Properties tab, modify the name of the policy or keep the default.
6. In the **Description** field, modify the description for the policy or keep the default.
7. In the **Re-arm Period** field, specify a re-arm period in minutes.

The rearm period is the amount of time after the policy executes before it can execute again. It is useful for limiting the number of times the same actions will execute.

IMPORTANT: Specify shorter periods for important applications.

8. In the **Condition** drop-down menu, specify a comparison operator.
9. (Available for only backup libraries). To the left of the comparison operator, select the media pool you want to monitor.
10. To the right of the comparison operator drop-down menu, do the following depending on the field displayed:
 - Type a percentage.
 - Type an amount in gigabytes (GB).
 - Type a number for available media. For example, if you want to be alerted when the number of available media for a storage pool is less than two, you would set the conditional to less than (<). You would then type 2 in the **Media** field.
11. Select an action to occur when the policy condition is fulfilled by clicking one of the following buttons (more than one action can be assigned to a policy):
 - **Send E-mail** - Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses, and then click **OK**.
 - **Generate Event** - Policy Manager generates an event of the specified event type. The event appears in Event Monitoring for Storage Essentials. Select one of the following event types.

IMPORTANT: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
- **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
- **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
- **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

After you select an event level, click **OK**.

- **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the %JBoss4_DIST%\server\appiq\remotescripts directory on the management server when the script is executed. You can use environment variables in your script, such as POLICY_NAME and POLICY_DESCRIPTION. POLICY_NAME provides the policy name and POLICY_DESCRIPTION provides the policy description. See ["Software Environment Variables for Scripting"](#) on page 207 for more information.

Prefix the command with "start" if the custom command triggers a user interface component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```

12. Click **Finish**.

13. To test a policy, click the **Test** button in the Utilization Policy table.

The management server fires a test for all utilization policies associated with that element.

Creating Policies for Discovery

You can create an infrastructure policy that generates an event, sends an e-mail, or runs a custom command when an element is discovered.

IMPORTANT: If you plan to use e-mail notification with your policy, first assign an SMTP server from which the management server can send its e-mail notifications. See ["Setting Up E-mail Notification"](#) on page 99.

To create a policy for discovery:

1. Access Policy Manager as described in the topic, ["Accessing Policy Manager"](#) on page 438.
2. In the Policy Manager tree in the middle pane, expand the node, Infrastructure Policies. Then, click **Provisioning**.
3. Click the **Add** button in the right pane.
4. In the **Name** field, type a name for the policy.
5. In the **Description** field, type a description for the policy.
6. Select one or more element types.
When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.
7. Select **Fire when event is cleared** if you want the policy to fire when the event is cleared, Otherwise the event fires when the event is received.
8. Select an action to occur when the policy condition is fulfilled by clicking one of the following buttons (more than one action can be assigned to a policy):

- **Send E-mail** - Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses, and then click **OK**.
- **Generate Event** - Policy Manager generates an event of the specified event type. The event appears in Event Monitoring for Storage Essentials. Select one of the following event types.

IMPORTANT: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
- **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
- **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
- **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

After you select an event level, click **OK**.

- **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the `%JBOS4_DIST%\server\appiq\remotescripts` directory on the management server when the script is executed. You can use environment variables in your script, such as `POLICY_NAME` and `POLICY_DESCRIPTION`. `POLICY_NAME` provides the policy name and `POLICY_DESCRIPTION` provides the policy description. See ["Software Environment Variables for Scripting"](#) on page 207 for more information.

Prefix the command with "start" if the custom command triggers a user interface component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```

9. Click **OK**.

Creating Policies for Provisioning

You can create an infrastructure policy that generates an event, sends an e-mail, or runs a custom command when successful provisioning occurred.

IMPORTANT: If you plan to use e-mail notification with your policy, first assign an SMTP server from which the management server can send its e-mail notifications. See "[Setting Up E-mail Notification](#)" on page 99.

To create a policy for provisioning:

1. Access Policy Manager as described in the topic, "[Accessing Policy Manager](#)" on page 438.
2. In the Policy Manager tree in the middle pane, expand the node, Infrastructure Policies. Then, click **Provisioning**.
3. Click the **Add** button.
4. In the **Name** field, type a name for the policy.
5. In the **Description** field, type a description for the policy.
6. Select one or more element types.

When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.
7. Select **Fire when event is cleared** if you want the policy to fire when the event is cleared, Otherwise the event fires when the event is received.
8. Select an action to occur when the policy condition is fulfilled by clicking one of the following buttons (more than one action can be assigned to a policy):
 - **Send E-mail** - Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses, and then click **OK**.
 - **Generate Event** - Policy Manager generates an event of the specified event type. The event appears in Event Monitoring for Storage Essentials. Select one of the following event types.

IMPORTANT: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
- **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
- **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
- **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

After you select an event level, click **OK**.

- **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Type a command that will execute the

script in the field. Then, click **OK**. The software assumes you are in the %JBOS4_DIST%\server\appiq\remotescripts directory on the management server when the script is executed. You can use environment variables in your script, such as POLICY_NAME and POLICY_DESCRIPTION. POLICY_NAME provides the policy name and POLICY_DESCRIPTION provides the policy description. See ["Setting Up E-mail Notification"](#) on page 99 for more information.

Prefix the command with "start" if the custom command triggers a user interface component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```

9. Click **OK**.

Creating Policies for Events

You can create a policy that generates an event, sends an e-mail, or runs a custom command when an event occurred on one or more specified elements

For example, you can create a policy that sends an e-mail when a new element generates a critical event.

If you plan to use e-mail notification with your policy, first assign an SMTP server from which the management server can send its e-mail notifications. See ["Setting Up E-mail Notification"](#) on page 99.

To create a policy for events:

1. Access Policy Manager as described in the topic, ["Accessing Policy Manager"](#) on page 438.
2. In the Policy Manager tree in the middle pane, expand the node, Infrastructure Policies. Then, click **Events**.
3. Click the **Add** button.
4. In the **Name** field, type a name for the policy.
5. In the **Description** field, type a description for the policy.
6. In the **Re-arm Period** field, specify a re-arm period in minutes.

The rearm period is the amount of time after the policy executes before it can execute again. It is useful for limiting the number of times the same actions will execute.

IMPORTANT: Specify shorter periods for important applications.

7. Select one or more element types.

When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.

8. Select **Fire when event is cleared** if you want the policy to fire when the event is cleared, Otherwise the event fires when the event is received.

9. In the **Severity** drop-down menu, specify a comparison operator.
10. To the right of the **Severity** drop-down menu, select one of the following severity levels:

IMPORTANT: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
 - **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
 - **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
 - **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
 - **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
 - **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.
11. Select one of the following from the **Summary Text** drop-down menu:
- **Is anything** - Regardless of the contents of the event's summary text, Policy Manager sends an e-mail, generates an event, or runs a custom command.
 - **Contains** - If the event's summary text contains the specified text, Policy Manager sends an e-mail, generates an event, or runs a custom command. Type the specified text in the field that appears to the right of the **Summary Text** drop-down menu.
 - **Matches Regular Expression** - If the event's summary text matches the specified expression, Policy Manager sends an e-mail, generates an event, or runs a custom command. Type the specified text in the field that appears to the right of the **Summary Text** drop-down menu.
12. Select an action to occur when the policy condition is fulfilled by clicking one of the following buttons (more than one action can be assigned to a policy):
- **Send E-mail** - Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses, and then click **OK**.
 - **Generate Event** - Policy Manager generates an event of the specified event type. The event appears in Event Monitoring for Storage Essentials. Select one of the following event types.

IMPORTANT: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
- **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
- **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.

- **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

After you select an event level, click **OK**.

- **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the %JBOS4_DIST%\server\appiq\remotescripts directory on the management server when the script is executed. You can use environment variables in your script, such as POLICY_NAME and POLICY_DESCRIPTION. POLICY_NAME provides the policy name and POLICY_DESCRIPTION provides the policy description. See ["Software Environment Variables for Scripting"](#) on page 207 for more information.

Prefix the command with "start" if the custom command triggers a user interface component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```

13. Click **OK**.

Testing a Utilization Policy

After you create or modify a utilization policy, test it to verify it provides the results you are anticipating. To test a policy, click the **Test** button in the Utilization Policy table. The management server fires a test for all utilization policies associated with that element.

Keep in mind the following:

- If you want to run the **Test** functionality more than once, set the Re-arm period to zero before clicking the **Test** button a second time.
- Policies that are triggered for virtual applications are also triggered for file servers.


Modifying Policies

This section describes the following:

- ["Modifying Utilization Policies"](#) on page 446
- ["Modifying Discovery Policies"](#) on page 448
- ["Modifying Provisioning Policies"](#) on page 449
- ["Modifying Policies for Events"](#) on page 450

Modifying Utilization Policies

To modify a utilization policy:

1. Access Policy Manager as described in the topic, "[Accessing Policy Manager](#)" on page 438.
2. Click the  button corresponding to the policy you want to modify.
3. In the **Name** field, change the name for the policy.
4. In the **Description** field, change the description for the policy.
5. Select or deselect the **Active** option.
When the **Active** option is selected, the policy is active.
6. In the **Re-arm Period** field, change the re-arm period in minutes.
The rearm period is the amount of time after the policy executes before it can execute again. It is useful for limiting the number of times the same actions will execute.

IMPORTANT: Specify shorter periods for important applications.

7. (Skip this step for backup clients) In the **Condition** drop-down menu, change the conditions of the policy.
8. To change the action to occur when the policy condition is fulfilled click one of the following buttons (more than one action can be assigned to a policy):
 - **Send E-mail** - Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses, and then click **OK**.
 - **Generate Event** - Policy Manager generates an event of the specified event type. The event appears in Event Monitoring for Storage Essentials. Select one of the following event types.

IMPORTANT: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
- **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
- **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
- **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

After you select an event level, click **OK**.

- **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the `%JBOS4_DIST%\server\appiq\remotescripts` directory on the management server when the script is executed. You can use environment variables in your script, such as

POLICY_NAME and POLICY_DESCRIPTION. POLICY_NAME provides the policy name and POLICY_DESCRIPTION provides the policy description. See ["Software Environment Variables for Scripting"](#) on page 207 for more information.

Prefix the command with "start" if the custom command triggers a user interface component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```

9. Click **OK**.

10. To test a policy, click the **Test** button in the Utilization Policy table.

The management server fires a test for all utilization policies associated with that element.

Modifying Discovery Policies

To modify a policy for discovery.

1. Access Policy Manager as described in the topic, ["Accessing Policy Manager"](#) on page 438.
2. In the Policy Manager tree in the middle pane, expand the node, Infrastructure Policies. Then, click **New Element Discovery**.

3. Click the  button corresponding to the policy you want to modify.

4. In the **Name** field, change the name for the policy.

5. In the **Description** field, change the description of the policy.

6. Select or deselect one or more element types.

When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.

7. Select **Fire when event is cleared** if you want the policy to fire when the event is cleared, Otherwise the event fires when the event is received.

8. Change the action to occur when the policy condition is fulfilled by clicking one of the following buttons (more than one action can be assigned to a policy):

- **Send E-mail** - Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses, and then click **OK**.
- **Generate Event** - Policy Manager generates an event of the specified event type. The event appears in Event Monitoring for Storage Essentials. Select one of the following event types.

IMPORTANT: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
- **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.

- **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
- **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

After you select an event level, click **OK**.

- **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the %JBOS4_DIST%\server\appiq\remotescripts directory on the management server when the script is executed. You can use environment variables in your script, such as POLICY_NAME and POLICY_DESCRIPTION. POLICY_NAME provides the policy name and POLICY_DESCRIPTION provides the policy description. See ["Software Environment Variables for Scripting"](#) on page 207 for more information.

Prefix the command with "start" if the custom command triggers a user interface component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```


9. Click **OK**.

Modifying Provisioning Policies

You can create an infrastructure policy that generates an event, sends an e-mail, or runs a custom command when successful provisioning occurred.

If you plan to use e-mail notification with your policy, first assign an SMTP server from which the management server can send its e-mail notifications. See ["Setting Up E-mail Notification"](#) on page 99.

To modify a policy for provisioning:

1. Access Policy Manager as described in the topic, ["Accessing Policy Manager"](#) on page 438.
2. In the Policy Manager tree in the middle pane, expand the node, Infrastructure Policies. Then, click **Provisioning**.
3. Click the  button corresponding to the policy you want to modify.
4. In the **Name** field, change the name for the policy.
5. In the **Description** field, change the description for the policy.
6. Select one or more element types.

When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.

7. Select **Fire when event is cleared** if you want the policy to fire when the event is cleared, Otherwise the event fires when the event is received.
8. Change the action to occur when the policy condition is fulfilled by clicking one of the following buttons (more than one action can be assigned to a policy):
 - **Send E-mail** - Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses, and then click **OK**.
 - **Generate Event** - Policy Manager generates an event of the specified event type. The event appears in Event Monitoring for Storage Essentials. Select one of the following event types.

IMPORTANT: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
- **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
- **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
- **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

After you select an event level, click **OK**.

- **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the %JBOS4_DIST%\server\appiq\remotescripts directory on the management server when the script is executed. You can use environment variables in your script, such as POLICY_NAME and POLICY_DESCRIPTION. POLICY_NAME provides the policy name and POLICY_DESCRIPTION provides the policy description. See ["Software Environment Variables for Scripting"](#) on page 207 for more information.

Prefix the command with "start" if the custom command triggers a user interface component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.


For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```

9. Click **OK**.

Modifying Policies for Events

To modify a policy for events:

1. Access Policy Manager as described in the topic, "[Accessing Policy Manager](#)" on page 438.
2. In the Policy Manager tree in the middle pane, expand the node, Infrastructure Policies. Then, click **Events**.
3. Click the  button corresponding to the policy you want to modify.
4. In the **Name** field, change the name for the policy.
5. In the **Description** field, type the description for the policy.
6. In the **Re-arm Period** field, change the re-arm period in minutes.

The rearm period is the amount of time after the policy executes before it can execute again. It is useful for limiting the number of times the same actions will execute.

IMPORTANT: Specify shorter periods for important applications.

7. Select or deselect one or more element types.
When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.
8. Select **Fire when event is cleared** if you want the policy to fire when the event is cleared, Otherwise the event fires when the event is received.
9. In the **Severity** drop-down menu, specify a comparison operator.
10. To the right of the **Severity** drop-down menu, you can change the severity level:

IMPORTANT: Since the severity level for an element is set by the manufacture, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
 - **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
 - **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
 - **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
 - **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
 - **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.
11. To change how Policy Manager scans the summary text, select one of the following from the **Summary Text** drop-down menu. Policy Manager scans the summary text and responds according to one of the following actions selected:

- **Is anything** - Regardless of the contents of the event's summary text, Policy Manager sends an e-mail, generates an event, or runs a custom command.
- **Contains** - If the event's summary text contains the specified text, Policy Manager sends an e-mail, generates an event, or runs a custom command. Type the specified text in the field that appears to the right of the **Summary Text** drop-down menu.
- **Matches Regular Expression** - If the event's summary text matches the specified expression, Policy Manager sends an e-mail, generates an event, or runs a custom command. Type the specified text in the field that appears to the right of the **Summary Text** drop-down menu.

12. Change the action to occur when the policy condition is fulfilled by clicking one of the following buttons (more than one action can be assigned to a policy):

- **Send E-mail** - Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses, and then click **OK**.
- **Generate Event** - Policy Manager generates an event of the specified event type. The event appears in Event Monitoring for Storage Essentials. Select one of the following event types.

IMPORTANT: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
- **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
- **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
- **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

After you select an event level, click **OK**.

- **Execute a Custom Command** - Policy Manager executes a custom command on the management server when the condition is fulfilled. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the `%JBOS4_DIST%\server\appiq\remotescripts` directory on the management server when the script is executed. You can use environment variables in your script, such as `POLICY_NAME` and `POLICY_DESCRIPTION`. `POLICY_NAME` provides the policy name and `POLICY_DESCRIPTION` provides the policy description. See ["Software Environment Variables for Scripting"](#) on page 207 for more information. Prefix the command with "start" if the custom command triggers a user interface component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```



13. Click **OK**.

Viewing Policies

To view policies, access Policy Manager as described in the topic, "[Accessing Policy Manager](#)" on page 438.

In the right pane, the policies are listed in a table.


Table 117 Policy Table Description

Heading	Description
Name	The name of the policy.
Description	A description of the policy.
Active	If there is a check mark in this column, the policy is active.
Edit	Click the  button to edit a policy, for example to change its Active status.
Delete	Click the  button to remove a policy.

Deactivating a Policy

Policies are activated when they are created. You can deactivate a policy, but still keep it stored in the management server. For example, assume you created a policy that sends an e-mail whenever an event of type minor is generated for a server. You might want to deactivate this policy before you upgrade the server.

To deactivate a policy:


1. Access Policy Manager as described in the topic, "[Accessing Policy Manager](#)" on page 438.
2. Click the  button corresponding to the policy you want to modify.
3. Deselect the **Active** option.
4. Click the **Finish** button.

The policy is deactivated.

Deleting Policies

To delete a policy:

1. Access Policy Manager as described in the topic, "[Accessing Policy Manager](#)" on page 438.

2. In the Policy Manager tree in the middle pane, click the element or infrastructure to view its policies.
3. Do one of the following:
 - **Delete a policy** - If you want to delete just one policy, click the  button corresponding to the policy you want to delete.
 - **Delete several policies at once** - If you want to delete several policies at once, select the check boxes next to the policies you want to delete. To select all of the policies, select the check box next to the **Name** heading. Then, click the **Delete** button.

The policies are deleted.

Providing E-mail Notification for a Policy


You can configure Policy Manager to provide e-mail notification when an element is being overused or when any of the following occurs:

- A new element is discovered
- Successful provisioning occurred
- An event occurred on one or more specified elements

Keep in mind the following:

- First assign an SMTP server from which the management server can send its e-mail notifications. See ["Setting Up E-mail Notification"](#) on page 99 for more information.
- The following instructions assume you have already created a policy. If you have not created a policy, see the following topics:
 - ["Creating a Utilization Policy"](#) on page 438
 - ["Creating Policies for Discovery"](#) on page 441
 - ["Creating Policies for Provisioning"](#) on page 442
 - ["Creating Policies for Events"](#) on page 444

To set up e-mail notification for a policy:

1. Access Policy Manager as described in the topic, ["Accessing Policy Manager"](#) on page 438.
2. Click the  button corresponding to the policy you want to modify.
3. Click the **Send E-mail** button.

Policy Manager sends an e-mail when the condition is fulfilled. The software verifies the address entered has a correct form.
4. Enter a comma-separated list of e-mail addresses, and then click **OK**.
5. Click the **Finish** button.

Providing Event Generation for a Policy

You can configure Policy Manager to generate an event when an element is being overused or when any of the following occurs:


- A new element is discovered

- Successful provisioning occurred
- An event occurred on one or more specified elements

The following instructions assume you have already created a policy. If you have not created a policy, see the following topics:

- ["Creating a Utilization Policy"](#) on page 438
- ["Creating Policies for Discovery"](#) on page 441
- ["Creating Policies for Provisioning"](#) on page 442
- ["Creating Policies for Events"](#) on page 444

To set up event generation for a policy:

1. Access Policy Manager as described in the topic, ["Accessing Policy Manager"](#) on page 438.
2. Click the  button corresponding to the policy you want to modify.
3. Click the **Generate Event** button.

Policy Manager generates an event of the specified event type. The event appears in Event Monitoring for Storage Essentials. Select one of the following event types.

IMPORTANT: Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** - The severity level is not known.
 - **Informational** - An example of an informational event is a progress report event for firmware download operation currently in progress.
 - **Warning** - An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
 - **Minor** - An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
 - **Major** - An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
 - **Critical** - An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.
4. After you select an event level, click **OK**.
 5. Click the **Finish** button.

Providing a Custom Command for a Policy


You can configure Policy Manager to run a custom command on the management server when an element is being overused or when any of the following occurs:

- A new element is discovered
- Successful provisioning occurred
- An event occurred on one or more specified elements

The following instructions assume you have already created a policy. If you have not created a policy, see the following topics:

- ["Creating a Utilization Policy"](#) on page 438
- ["Creating Policies for Discovery"](#) on page 441
- ["Creating Policies for Provisioning"](#) on page 442
- ["Creating Policies for Events"](#) on page 444

To set up a custom script for a policy:

1. Access Policy Manager as described in ["Accessing Policy Manager"](#) on page 438.
2. Click the  button corresponding to the policy you want to modify.
3. Click the **Execute Custom Command** button.
4. Type a command that will execute the script in the field. Then, click **OK**. The software assumes you are in the %JBOS4_DIST%\server\appiq\remotescripts directory on the management server when the script is executed. You can use environment variables in your script, such as POLICY_NAME and POLICY_DESCRIPTION. POLICY_NAME provides the policy name and POLICY_DESCRIPTION provides the policy description. See ["Software Environment Variables for Scripting"](#) on page 207 for more information.

Prefix the command with "start" if the custom command triggers a user component, such as a program that uses the user interface (Example: Microsoft Internet Explorer) or a command prompt window.

For example, assume you want the custom command to open a command prompt window and list the contents in the directory. You would prefix the command with start as shown below:

```
start dir
```

5. Click the **Finish** button.

Policy Manager executes a remote script on the management server when the condition is fulfilled.

16 Chargeback Manager

IMPORTANT: Depending on your license, Chargeback Manager may not be available. See the “List of Features” to determine if you have access to Chargeback Manager. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).

This chapter describes the following:

- “[About Chargeback Manager](#)” on page 457
- “[Setting Up Chargeback Manager](#)” on page 458
- “[Accessing Chargeback Manager](#)” on page 459
- “[Creating an Asset Record](#)” on page 459
- “[Changing the Status of an Element](#)” on page 460
- “[Saving Chargeback Manager Information](#)” on page 461
- “[Viewing Assets](#)” on page 461
- “[Setting Storage Tier Name and Cost](#)” on page 462
- “[Setting the Infrastructure Cost](#)” on page 463
- “[Adding Asset Information](#)” on page 463
- “[Managing Departments](#)” on page 466
- “[Setting Up Asset and Storage Based Chargeback Manager](#)” on page 468
- “[Viewing Chargeback](#)” on page 481
- “[Chargeback Reports](#)” on page 484
- “[Filtering Assets](#)” on page 491

About Chargeback Manager

Chargeback Manager helps you manage departmental ownership, track cost and assemble business reports making inquiries, such as audits and inventory reviews, easier.

The management server provides the following types of chargeback.

- **Asset-based** - Asset-based chargeback calculates chargeback based on the departmental ownership percentages and the depreciated value of the assets. Each piece of equipment is owned by a department or a set of departments. Each department has a percentage ownership of the equipment.
- **Storage-based** - Storage-based chargeback calculates charges based on the actual amount of storage used by an application, the type of storage it is using and the ownership percentage assigned to each department. The chargeback number is further refined by an additional fixed infrastructure tax on a per department basis.

After you add information about all of your assets, back up the database by using the Database Admin Utility. Backing up the database saves your chargeback information. If the database fails,

your asset information is restored when you restore the database. See [“Performing an RMAN Hot Backup”](#) on page 148 for more information.

First set up your chargeback as described in the topic, [“Setting Up Chargeback Manager”](#) on page 458. When you are done with adding your chargeback information, you can view chargeback as following:

- **By element** - Displays chargeback for a single element. See [“Viewing Chargeback by Element”](#) on page 481 for more information.
- **By department** - Displays chargeback for a department. See [“Viewing Chargeback by Department”](#) on page 482 for more information.
- **By owner** - Displays chargeback for by owner. See [“Viewing Chargeback by Owner”](#) on page 483 for more information.

Chargeback Manager also provides the following reports. See [“Viewing Chargeback Reports”](#) on page 484 for more information.

- **Array Based Assets** - Displays asset information from storage arrays: host name, department, HBA port, HBA port WWN, storage volume, volume size, and cost.
- **Asset Based** - Displays the following asset-based chargeback for each department owning elements: department, asset name, ownership ratio, and chargeback amount. Total asset-based cost per month is also displayed.
- **Storage Based** - Displays storage-based chargeback for each department owning applications: department, application, ownership ratio, total capacity, and chargeback amount. Total storage-based cost per month is also displayed.
- **Storage System Based by Tier** - Displays storage-based chargeback by tier. Only storage systems are assigned storage tiers.

After viewing the reports, you can e-mail them to co-workers. You can even set up a schedule for a chargeback report to be mailed at regular intervals to a co-worker. See the topics, [“E-mailing a Chargeback Report”](#) on page 485 and [“Adding an E-mail Schedule for a Chargeback Report”](#) on page 486 for more information.

Chargeback Manager helps you track of the status of your elements. Elements that have recently been discovered are automatically given the status of New. You can then change the status of elements to in use, missing or repaired. Since the management server cannot determine what you plan to do with an element, you must change the status manually. However, you can easily change the status of a group of elements at once. See [“Changing the Status of an Element”](#) on page 460 for more information.

Chargeback Manager also provides a brief listing of your assets by name, status, serial number, vendor/model, and description. You can filter elements by status and/or type for easy navigation. You can even create custom filters. See [“About Filtering Assets”](#) on page 492.

Setting Up Chargeback Manager

To be able to view chargeback, you must first complete the steps in the following table.

Table 118 Setting Up Chargeback Manager

Step	Description	Where to Find More Info
1	Create an asset record if it does not exist. You can use Chargeback Manager for applications and hardware the management server does not detect. Just create an asset record for that element.	" Creating an Asset Record " on page 459
2	Set the status of the asset.	" Changing the Status of an Element " on page 460
3	(Optional) Add asset information for asset management.	" Adding Asset Information " on page 463
4	Add departments.	" Adding Departments " on page 467
5	Set up Chargeback Manager.	<ul style="list-style-type: none">• "Setting Up Asset-Based Chargeback Manager" on page 468• "Setting Up Storage-Based Chargeback Manager" on page 471

Accessing Chargeback Manager


To access Chargeback Manager, click **Tools > Storage Essentials > Chargeback Manager**.

Creating an Asset Record

IMPORTANT: Only a user belonging to a role that has System Configuration selected on the Edit Role page (such as the Domain Administrator role) is allowed to create a record.

You can use Chargeback Manager for any application or hardware, even those the management server does not detect. Just create an asset record for that application or hardware the management server does not monitor. Then, follow the steps for setting up Chargeback Manager, as described in the topic, "[Setting Up Chargeback Manager](#)" on page 458.

After you create a record, the element you created the record for is treated as a discovered element. A discovered element is an element that has been detected by the management server, but the management server cannot obtain detailed information about the element. If you create a record for an application, that application is treated as a virtual application.

IMPORTANT: You can easily remove an element's record by clicking the  button. When you remove an element's record, the management server no longer monitors that element. See ["Deleting Elements from the Management Server"](#) on page 38 for more information.

To create a record:

1. Access Chargeback Manager, as described in ["Accessing Chargeback Manager"](#) on page 459.
2. At the top of the right pane, click the **New** button.
3. In the **Create Asset Record** window, type the following:
 - Name
 - Vendor
 - Model
4. Select one of the following for the type of element:
 - **Hardware** - Host
 - **Hardware** - Storage System
 - **Hardware** - Switch
 - **Software** - Application
5. Click **OK**.

Changing the Status of an Element

Chargeback Manager helps you track of the status of your elements. Elements that have recently been discovered are automatically given the status of New. You can then change the status of elements to in use, missing or repaired. Since the management server cannot determine what you plan to do with an element, you must change the status manually. However, you can easily change the status of a group of elements at once.

IMPORTANT: Once you are done with change the status of your elements, save your settings. See ["Saving Chargeback Manager Information"](#) on page 461.

To change the status of an element:

1. Access Chargeback Manager, as described in ["Accessing Chargeback Manager"](#) on page 459.
2. Under the Status column, select the status of the element:
 - **New (Default)** - The status of the element has not been set yet.
 - **Missing** - The element cannot be found. It may have been taken off line.
 - **Repaired** - The element is repaired.
 - **In Use** - The element is running.

The status settings are set manually. For example, if the status of an element changes from being “In Use” to being “Repaired,” you must change this status manually. Refer to the [Topology and Event Monitoring for Storage Essentials](#) for the latest status of an element.

3. To change the status of multiple elements at once:
 - a. Select the elements you want to modify.
 - b. Click the **Set Status** button.
 - c. Select the new status for the elements you selected from the Asset Status drop-down menu.
 - d. Click **OK**.

Saving Chargeback Manager Information

After you change the status of your elements, save your settings by clicking the **Save Listing to File** link. The following information is saved as comma-separated values, which can be viewed using a text editor, such as Notepad, or a spreadsheet program, such as Microsoft Excel.

- **ID**
- **Name**
- **Status**
- **Category**
- **Serial Number**
- **Vendor**
- **Model**

Viewing Assets

To obtain asset information about an element:

1. Access Chargeback Manager, as described in [“Accessing Chargeback Manager”](#) on page 459.

The following is displayed in the right pane:







- **Name** - The name of the element.
- **Status** - The state of the element. An element is automatically assigned the status of “new” when it is first discovered. You can change the status of an element to in use, repaired or missing. See [“Changing the Status of an Element”](#) on page 460 for more information.
- **Chargeback Manager** - Click the  icon to view chargeback for an element. You must first set up Chargeback Manager before you can view its calculations. See [“Setting Up Chargeback Manager”](#) on page 458 for more information.
- **Vendor/Model** - The vendor and/or model of an element.
- **Type** - The type of element, such as an application. See [Table 119](#) on page 462 for an explanation of each icon.
- **Serial Number** - The serial number of the element.
- **Description** - Description of the element.

Table 119 Element Type Icons

Graphic	Element Type
	Application
	Host
	Switch
	Storage System

2. To remove an asset record, click the  button corresponding to the record you want to remove.

Setting Storage Tier Name and Cost

Storage-based chargeback lets you charge the application owners based on the amount of storage allocated to them. Each storage system is designated a storage tier classification. You can specify a charge for each storage tier. This charge is referred to as storage tier cost. The management server determines the storage cost of the application by multiplying the storage tier cost by the allotted storage. If the application uses more than one storage system, the storage cost from each storage system is added for the total storage cost.

Storage tiers for storage-based information can have any name. The following storage tiers names are provided.

- **No Tier** is the default setting. The storage is not classified.
- **Ultra High Availability** is usually assigned to the ultra high availability storage. This tier contains the premium storage in your network, usually the most expensive.
- **High Availability** is usually assigned to the high availability storage. This tier contains storage that is not as expensive as the storage assigned to the Ultra High Availability tier.
- **IDE Based Storage** is usually assigned to IDE-based storage. This tier contains storage that is comparatively inexpensive.

Do the following:

1. Access Chargeback Manager, as described in "[Accessing Chargeback Manager](#)" on page 459.
2. Click the **Ownership** tab in the right pane.
3. Click the **Set Storage Tier Cost** button.
4. Add the cost for each tier.
5. To change the name of a storage tier, type a new name for the storage tier in the **Storage Tier Name** field.

Keep in mind the following:

- The name must contain 1 to 256 characters.

- The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or one of the following symbols: dollar sign (\$), caret (^), hyphen (-), an underscore (_) or a space.

6. When you are done, click **OK**.

Setting the Infrastructure Cost

Specify a monthly infrastructure charge for when storage-based calculation is being done. This charge is identical for each department and is applied once each month on top of their total storage cost. Modifying the infrastructure charge will impact the storage-based chargeback result for all department owners. The same infrastructure charge is assigned to a department regardless of it owning one or 100 elements.

For example, the infrastructure cost is not included when you view the chargeback for individual elements; however, the infrastructure cost is added to the Total Cost/Month value in the storage-based Chargeback Manager report. It is also displayed when you view chargeback per department. The infrastructure cost is added to each department, regardless of whether the department owns one or 100 elements.

1. Access Chargeback Manager, as described in "[Accessing Chargeback Manager](#)" on page 459.
2. Click the **Ownership** tab in the right pane.
3. Click the **Set Infrastructure Cost** button.
4. Type the monthly infrastructure charge for when asset-based calculation is performed.
5. Click **OK**.

Adding Asset Information

This section describes the following:


- "[Adding Asset Information](#)" on page 463
- "[Adding General Information](#)" on page 464
- "[Adding Staff Information](#)" on page 465
- "[Adding Geographic Information](#)" on page 466
- "[Adding Licensing and Warranty Information](#)" on page 466
- "[Adding Custom Information](#)" on page 466

Adding Asset Information

Chargeback Manager provides a handy way for you to keep track of your asset information for an element. You can easily store warranty and licensing information, as well as contact information for the element. For example, assume a switch on the network is having some problems, and you want to contact the person in charge of that switch. You can use the element's asset record to not only find the contact information for that switch, but also the location of that switch.

IMPORTANT: After you add information about all of your assets, back up the database. Backing up the database saves your chargeback information. If the database fails, your asset information is restored when you restore the database. See ["Performing an RMAN Hot Backup"](#) on page 148 for more information.

You can view and add asset information by doing the following:

1. Access Chargeback Manager, as described in ["Accessing Chargeback Manager"](#) on page 459.
2. Do one of the following:
 - Click the  icon corresponding to the element.
 - Click the link for the element in the right pane.
3. To access the different types of asset information, click the Asset Record node or the nodes under it. To view general information about an element, click the Asset Record node in the tree. To view specific asset information such as ownership, click the Staff node.

You can also access the tree from Application Viewer and System Manager:

- To access the tree from Application Viewer, click the name of an application in the Application Viewer tree. In the right pane, click the **Asset Management** tab.
- To access the tree from System Manager, double-click an element in the topology. In the right pane, click the **Asset Management** tab.

NOTE: The nodes under Chargeback Manager in the tree are for creating and viewing reports. The Storage-based node is available only for applications. See ["Chargeback Reports"](#) on page 484 for more information.

Adding General Information

The management server provides a page for you to enter the following general information about an element. When you are done with adding information on this page, click the **Save Changes** button at the bottom of the page. To learn more about Chargeback Manager, see ["About Chargeback Manager"](#) on page 457. and ["Setting Up Chargeback Manager"](#) on page 458.

NOTE: The fields that accept input cannot contain more than 250 characters.

- **Custom Name** - A name you assigned to the element. See ["Assigning a Custom Name"](#) on page 221 for more information.
- **Date Created** - Date the element was discovered.
- **Date Last Modified** - Date the record was last modified.
- **Description** - A description of the element. This description cannot be more than 250 characters.

- **Status** - The current status of the element. If the status of the element has changed, select the new status from the Status drop-down menu.
 - **New** - Default category for all detected elements.
 - **Missing** - No longer detectable through discovery
 - **Repaired** - The element is being repaired. The software does not automatically select this status.
 - **In Use** - The element is in use.
- **Vendor** - The vendor for the element.
- **Model** - The model of the element.
- **Serial Number** - Serial number of the element.
- **Barcode Number** - The barcode on the device.
- **Asset Code** - The asset code assigned to the element.
- **Asset Type** - The asset type assigned to the element.
- **Asset Tag** - The asset tag assigned to the element.
- **Asset Category** - The asset category assigned to the element.
- **Geographic Location** - The location of the element, for example, Boston, Massachusetts.
- **(Storage Systems Only) Storage Tier Classification** - Select the storage tier you want to assign the storage system. Storage tiers for storage-based information can have any name. By default, they are the following:
 - **No Tier** is the default setting. The storage is not classified.
 - **Ultra High Availability** is usually assigned to the ultra high availability storage. This tier contains the premium storage in your network, usually the most expensive.
 - **High Availability** is usually assigned to the high availability storage. This tier contains storage that is not as expensive as the storage assigned to the Ultra High Availability tier.
 - **IDE Based Storage** is usually assigned to IDE-based storage. This tier contains storage that is comparatively inexpensive.

To set the cost assigned to a storage tier, click the **Set Storage Tier Cost** link. Enter the cost for each tier. Then, click **OK**. See "[Setting Storage Tier Name and Cost](#)" on page 462 for more information about changing storage tier names and adding cost information.

Adding Staff Information

This page provides contact information about the element.

IMPORTANT: Keep this information up to date. Other users will use this information to contact you about the element, for example, if it is having problems.

- **Administrator** - The person or department that maintains the element.
- **Staff Name** - The name of the person who maintains the element.
- **Staff Phone Number** - A phone number for the person who maintains the element
- **Staff Department** - The department that maintains the element

- **Staff E-Mail** - An e-mail address of the person who maintains the element

Adding Geographic Information

Use this page to add geographic information about the element. This page is helpful in keeping track where all of your elements are located, especially if you have more than 100 elements. For example, assume you are told one of your servers is having problems and you need physical access to it. You can use this page to find where the server is located.

- **Rack Number** - The number of the rack that holds the element.
- **Floor** - The floor on which the element is located, for example third floor.
- **Data Center** - The name of the data center where the element is located.
- **Address** - The address where the element is located.
- **City** - The city where the element is located, for example, Boston.
- **Region** - The region where the element is located, for example, New England.
- **Country** - The country where the element is located, for example, the United States.
- **Continent** - The continent where the element is located, for example, North America.
- **Zip Code** - The zip code for the town where the element is located. For example, if the element is located in Burlington, Massachusetts, the zip code would be 01803. If your country does not use zip codes, you can leave this field blank.

Adding Licensing and Warranty Information

Use this page to provide licensing and warranty information.

- **License (maximum of 4000 characters)** - The license of the element.
- **Warranty Information (maximum of 4096 characters)** - Information about the warranty. In this field, you probably want to enter information such as how long the warranty lasts and what it covers.
- **Comments (maximum of 4000 characters)** - Any additional financial information you might want to add about the element

Adding Custom Information

You can provide up to six custom properties.

1. In the **Name** field, assign a name for the field. The name cannot be more than 50 characters. For example, Backup Contact.
2. In the **Value** field, provide the information for the field. Do not provide more than 255 characters. For example, Joe Smith.
3. Repeat the previous step for each custom property you want to add.
4. When you are done, click **Save Changes**.

Managing Departments

This section describes the following:

- ["Adding Departments"](#) on page 467

- ["Editing a Department"](#) on page 467
- ["Removing a Department from Chargeback Manager"](#) on page 467

Adding Departments

Before you can assign a department to an element, you must add it to the list of departments, as described in the following steps.

To add a department:


1. Access Chargeback Manager, as described in ["Accessing Chargeback Manager"](#) on page 459.
2. Click the **Departments** tab above the table.
3. Click the **New** button.
4. In the **Add Department** window, provide the following information:
 - Department Name (Required)
 - Department Number (Required)
 - E-mail
 - Phone

5. Click **OK**.

The new department is added.

Editing a Department

To edit a department:

1. Access Chargeback Manager, as described in ["Accessing Chargeback Manager"](#) on page 459.
2. Click the **Departments** tab above the table.
3. Click the  button corresponding with the department you want to edit.
4. In the **Edit Department** window, you can edit all fields except the department number.
5. Click **OK**.


Removing a Department from Chargeback Manager

Over time, departments in your company may merge and others may be dissolved. To keep up with these changes, you may need to remove obsolete departments from your list. If an element is only assigned to this department, it no longer has an owner. Likewise, if an element is assigned to this department and several others, it continues to be assigned to the other departments.

For example, assume you want to delete a department called TooSmall. The TooSmall department owns 50 percent of a host and the Server department owns 50 percent of the host. When you remove TooSmall, the host is only owned by the Server department by 50 percent.

To remove a department from the list:

1. Access Chargeback Manager, as described in ["Accessing Chargeback Manager"](#) on page 459.

2. Click the **Departments** tab above the table.
3. Click the **Delete** button () corresponding to the department you want to remove.

Setting Up Asset and Storage Based Chargeback Manager

This section describes the following:

- ["Setting Up Asset-Based Chargeback Manager"](#) on page 468
- ["Setting Up Storage-Based Chargeback Manager"](#) on page 471
- ["Editing Percentage of Ownership"](#) on page 475
- ["Removing Department Ownership of an Element"](#) on page 475
- ["How a Depreciation Method Is Calculated"](#) on page 476
- ["Capacity in Chargeback Manager and Capacity Manager Differ"](#) on page 476

Setting Up Asset-Based Chargeback Manager

Asset-based chargeback calculates chargeback based on the departmental ownership percentages and the depreciated value of the assets. Each piece of equipment is owned by a department or a set of departments. Each department has a percentage ownership of the equipment.


The management server calculates monthly chargeback from the financial information provided. You can then use these monthly calculations to determine the cost impact on your enterprise on a monthly basis. You can even break the cost down by department. If you have a infrastructure cost, you can add that into the calculations as well.

To obtain this asset-based chargeback, you must follow these steps:

- Step 1 - Specify Financial information.
- Step 2 - Assign Departmental Ownership Percentage.
- Step 3 - Review asset-based chargeback result.

IMPORTANT: You must have already added your departments, as described in the topic, ["Editing a Department"](#) on page 467.


To provide information for asset-based chargeback:

1. Access Chargeback Manager, as described in ["Accessing Chargeback Manager"](#) on page 459.
2. Do one of the following:
 - Click the  icon corresponding to the element.
 - Click the link for an element in the right pane.
3. Click the **Asset-based** node under the Chargeback Manager node.

NOTE: You can also access the tree shown in this figure from Application Viewer and System Manager:

- To access the tree from Application Viewer, click the name of an application in the Application Viewer tree. In the right pane, click the **Asset Management** tab.
- To access the tree from System Manager, double-click an element in the topology. In the right pane, click the **Asset Management** tab.

Step 1 - Specify Financial information

1. Verify the option, **Step 1 - Specify Financial information**, is selected at the top of the page.
2. Provide the following financial information:
 - **Purchase Order Number** - The purchase order of the element
 - ***Date Purchased** - The date the element was purchased. To select the date, Click the calendar icon, . If you select a future date, the purchase date is set to today when calculating depreciation. The management server only supports dates within the years 1900 and 3000.
 - **Reseller** - The company that directly sold you the element.
 - ***Purchase Price** - The price of the element when it was bought
 - **Salvage Value** - The amount of money an item is worth for salvage value. You cannot go below this number when depreciating an item.
 - ***Depreciation Period (months)** - This time period in which you plan to keep the element.
 - ***Depreciation Method** - Select how the depreciation is calculated:
 - **Straight Line** - The device loses the same amount of value in each period. To learn more about how the management server calculates straight line depreciation, see "[Calculating Straight Line Depreciation](#)" on page 477.
 - **Fixed Declining Balance** - The Fixed Declining Balance method calculates depreciation based on the value of the asset each month instead of a fixed rate like straight line depreciation. To learn more about how the management server calculates fixed declining balance, see "[Calculating Fixed Declining Balance](#)" on page 478.
 - **Double Declining Balance** - This depreciation method doubles the calculation of the Fixed Declining Balance method. Thus, it doubles the speed at which a device depreciates. To learn more about how the management server calculates double declining balance, see "[Calculating Double Declining Balance](#)" on page 479.
 - **Value** - The value of the element as of the end of the previous month. For example, assume you entered and/or viewed this chargeback information in the middle of March. The value would be for the month of February, but not for March. The value is calculated from the following fields:
 - **Date Purchased**
 - **Original Cost (\$)**
 - **Depreciation Salvage Value (\$)**

- **Depreciation Period**
- **Depreciation Method**

*Required

3. Click **Save Changes**.

Step 2 - Assign Departmental Ownership Percentage

Assign the percentage of ownership to an element and the monthly infrastructure charge by doing the following:

1. Select the option, **Step 2 - Assign Departmental Ownership Percentage.**, at the top of the page.
2. Click the **Add Ownership** button.
3. Select a department from the **Department** drop-down menu.
4. Type the percentage of ownership in the **Ownership %** field. If you do not see a department listed, add it to the list as described in the topic, "[Adding Departments](#)" on page 467. Click the **Manage Departments** link. After you have added the department, close the window you used to add the department and then refresh the page.
5. Click **OK**.

The department with its percentage of ownership is added to the table.

6. If multiple departments own the element repeat the steps above for each element. You can have departments owning more than 100 percent of the element.
7. (Optional) Specify a monthly infrastructure charge for when asset-based calculation is being done. This charge is identical for each department and is applied once each month on top of their total ownership cost. Modifying the infrastructure charge impacts the asset-based chargeback result for all department owners.

For example, the infrastructure cost is not included when you view the chargeback for individual elements; however, the infrastructure cost is added to the Total Cost/Month value in the Asset-based Chargeback report. It is also displayed when you view chargeback per department. The infrastructure cost is added to each department, regardless of whether the department owns one or 100 elements.

- a. Click the **Set Infrastructure Cost** button.
 - b. Type the monthly infrastructure charge for when asset-based calculation is performed.
 - c. Click **OK**.
8. When you are done with assigning the element to a department, click the **Save Changes** button.

Step 3 - Review Asset-based Chargeback Result

IMPORTANT: The management server displays chargeback information up to the end of the previous month. For example, assume you view chargeback information in the middle of March. The calculations for chargeback would include the month of February, but not March.

To view the result of Asset-based Chargeback:

1. Select the option, **Step 3 - Review Asset-based Chargeback Result.**, at the top of the page.
2. If you see empty values, verify you have provided the required values in the previous steps.

The ownership cost is determined by the following formula:

$$(\text{Depreciation}) \times (\text{Ownership \%}) = \text{Ownership Cost}$$

Ownership Cost is how much owning the element will cost a department. The depreciation is determined by the depreciation method you selected in Step 1 - Specify Financial information.

IMPORTANT: The infrastructure cost is not included in ownership cost because the information displayed on this page is per asset. The asset-based infrastructure cost is a monthly charge that is applied to each departmental owner in addition to any ownership charges. The infrastructure cost is not included when you view the chargeback for individual elements; however, the infrastructure cost is added to the Total Cost/Month value in the Asset-based Chargeback report. It is also displayed when you view chargeback per department. The infrastructure cost is added to each department, regardless of whether the department owns one or 100 elements.

Setting Up Storage-Based Chargeback Manager

Storage-based Chargeback Manager calculates charges based on the actual amount of storage used by an application on the storage system, the type of storage it is using and the ownership percentage assigned to each department. The chargeback number is further refined by an additional fixed infrastructure tax on a per department basis.

To obtain this storage-based chargeback, you must follow these steps:


- Step 1 - Assign Departmental Ownership Percentage.
- Step 2 - Review Storage Tier Cost.
- Step 3 - Review Storage Dependency and Cost.
- Step 4 - Review storage-based chargeback result.

Keep in mind the following:

- You must have already added your departments, as described in the topic, ["Adding Departments"](#) on page 467.
- Not all applications use storage on storage systems. Storage-based chargeback is applicable only for those applications that use storage on storage systems.
- You must have access to the storage system the application uses. Verify your organization and roles allow you access. See ["About the Security for the Management Server"](#) on page 71 for more information.
- Chargeback Manager provides only network capacity. If you look at the capacity of an application in Capacity Manager, the capacity differs. Capacity Manager displays the total capacity of an application, including the network drives. See the topic, ["Capacity in Chargeback Manager and Capacity Manager Differ"](#) on page 476 for more information.

To provide information for storage-based Chargeback Manager:

1. Access the management server.

2. Click **Chargeback Manager** in the left pane.
3. Do one of the following:
 - Click the  icon corresponding to the element.
 - Click the link for a host running an application in the right pane.
4. Click the **Storage-based** node under the Chargeback Manager node.

NOTE: You can also access the tree shown in this figure from Application Viewer and System Manager.

- To access the tree from Application Viewer, click the name of an application in the Application Viewer tree. In the right pane, click the **Asset Management** tab.
- To access the tree from System Manager, double-click an element in the topology. In the right pane, click the **Asset Management** tab.

Step 1 - Assign Departmental Ownership Percentage

1. Select the option, **Step 1 - Assign Departmental Ownership Percentage.**, at the top of the page.
2. Click the **Add Ownership** button.
3. Select department from the **Department** combo-box.

NOTE: If you do not see a department listed, add it to the list as described in the topic, "[Adding Departments](#)" on page 467. Click the **Manage Departments** link. After you have added the department, close the window you used to add the department and then refresh the page.

4. Type the percentage of ownership in the **Ownership %** field.
5. Click **OK**.

The department with its percentage of ownership is added to the table.
6. If multiple departments own the element repeat the steps above for each element. You can have departments owning more than 100 percent of the element.
7. (Optional) Specify a monthly infrastructure charge for when storage-based calculation is being done. This charge is identical for each department and is applied once each month on top of their total storage cost. Modifying the infrastructure charge will impact the storage-based chargeback result for all department owners. The same infrastructure charge is assigned to a department regardless of it owning one or 100 elements.

For example, the infrastructure cost is not included when you view the chargeback for individual elements; however, the infrastructure cost is added to the Total Cost/Month value in the Storage-based Chargeback report. It is also displayed when you view chargeback per department. The infrastructure cost is added to each department, regardless of whether the department owns one or 100 elements.

 - a. Click the **Set Infrastructure Cost** button.
 - b. Type the monthly infrastructure charge for when asset-based calculation is performed.

- c. Click **OK**.
8. When you are done with assigning the element to a department, click the **Save Changes** button.

Step 2 - Review Storage Tier Cost

Storage-based chargeback lets you charge the application owners based on the amount of storage allocated to them. Each storage system is designated a storage tier classification. You can specify a charge for each storage tier. This charge is referred to as storage tier cost. The management server determines the storage cost of the application by multiplying the storage tier cost by the allotted storage. If the application uses more than one storage system, the storage cost from each storage system is added for the total storage cost.

Storage tiers for storage-based informations can have any name. By default, they are the following:

- **No Tier** - the default setting. The storage is not classified.
- **Ultra High Availability** - usually assigned to the ultra high availability storage. This tier contains the premium storage in your network, usually the most expensive.
- **High Availability** - usually assigned to the high availability storage. This tier contains storage that is not as expensive as the storage assigned to the Ultra High Availability tier.
- **IDE Based Storage** - usually assigned to IDE-based storage. This tier contains storage that is comparatively inexpensive.

Do the following:

1. Add the cost for each tier.
2. (Optional) To change the name of a storage tier, type a new name for the storage tier in the Storage Tier Name field.

Keep in mind the following:

- The name must contain 1 to 256 characters.
 - The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or one of the following symbols: dollar sign (\$), caret (^), hyphen (-), an underscore (_) or a space.
3. When you are done, click **Save Changes**.

Step 3 - Review Storage Dependency and Cost

IMPORTANT: The management server displays chargeback information up to the end of the previous month. For example, assume you view chargeback information in the middle of February. The calculations for chargeback would include the month of January, but not February.

The Storage Dependency and Cost table is displayed if all of the following conditions apply:

- The application depends on a storage system.
- The organizations to which you belong allow you to view the storage system.

NOTE: You cannot change the storage tier if your role does not allow you to make modifications to the storage system.

If the table is empty and you know the application is dependent on a storage system, verify you have access to the storage system; otherwise, data cannot be calculated for this report.

If the application depends on a storage system, such as the Microsoft Exchange application in the following figure, the management server determines the storage cost by adding the mounted and unmounted storage, and multiplying the total storage storage in gigabytes by the tier cost per gigabyte:

$$(\text{Storage in gigabytes}) \times (\text{Tier Cost per gigabyte}) = \text{Storage Cost}$$

The Tier Cost was determined in Step 2 - Review Storage Tier Cost. Select a storage tier and then click **Save Changes**. The value for allotted storage includes mounted and unmounted storage. Any volumes the application can access are included in the storage calculations.

Keep in mind the following:

- The Tier Cost per GB and Storage Cost columns are populated after you click the **Save Changes** button.
- The name of the tier classifications shown in the figure can be customized, and therefore they can differ from the names shown in the figure in this section.

IMPORTANT: If you perform the calculations manually, you may see a different value than the one displayed in the Storage Cost column. When the management server calculates the storage cost, it does not use the rounded values displayed in the Mounted Storage, Unmounted Storage and Storage Tier columns. The management server rounds the values of all calculation results, except for the total storage cost (Step 4), in the tables for better viewing. In the case of US currency, it rounds the values off to the penny. To preserve accuracy, the management server does not use the rounded values displayed in the tables when it performs any of its calculations.

Step 4 - Review Storage-Based Chargeback Result

IMPORTANT: The management server displays chargeback information up to the end of the previous month. For example, assume you view chargeback information in the middle of February. The calculations for chargeback would include the month of January, but not February.

The total storage-based monthly chargeback is displayed. This number is calculated as follows:

$$(\text{Total Storage Cost}) \times (\text{Ownership \%}) = \text{Ownership Cost}$$

where


- **Total Storage Cost** is the total Storage Cost from (Step 3 - Review Storage Dependency and Cost).
- **Ownership %** is the percentage of ownership.

- The following figure displays the Chargeback result. To find the total storage cost, the management server added the values in the Storage Cost column in the previous figure (\$458.08 + \$698.62). The management server then multiplied the percentage of ownership, which in this case was 100 percent for the Engineering department. The resulting ownership cost is displayed below.

IMPORTANT: You may have noticed a discrepancy in the following figure. When you look at the Total Storage Cost value in the table, you see many numbers after the decimal place. The management server lets you see the true value of the number without rounding. The management server does round the value for the total storage cost above the table, so that you can provide a realistic figure to upper-management. In the case of US currency, the total storage cost is rounded off to the penny.

Editing Percentage of Ownership


To remove the department ownership of an asset:

1. Access the management server.
2. Click **Chargeback** in the left pane.
3. Click the link for the element in the right pane.
4. Click Asset-based under the Chargeback node in the tree.
5. Verify the option, "Step 2 - Assign Departmental Ownership Percentage" is displayed in the right pane..
6. Click the Edit () button corresponding to the percentage of ownership you want to modify.
7. In the **Ownership %** field, type a new percentage of ownership.
8. Click **Save Changes**.

Removing Department Ownership of an Element

Sometimes you may need to remove ownership from an element, such in the case of an element being moved from one department to another. When department ownership is removed from an element, the department is still accessible from the list of departments. If you want to make the department inaccessible to all elements, remove from the list of departments as described in the topic, "[Removing a Department from Chargeback Manager](#)" on page 467.

To remove ownership:

1. Access the management server.
2. Click Chargeback Manager in the left pane.
3. Click the link for the element in the right pane.
4. Click Asset-based under the Chargeback node in the tree.
5. Verify the option, "Step 2 - Assign Departmental Ownership Percentage" is displayed in the right pane..
6. Click the Delete () button corresponding to the department you want to remove.

7. Click **Save Changes**.

The department is removed.

Capacity in Chargeback Manager and Capacity Manager Differ

The capacity displayed for an application in Chargeback Manager differs from the capacity displayed in Capacity Manager. The management server uses only network storage when calculating chargeback capacity. Local capacity is not counted. The following figure shows the chargeback capacity for an Oracle instance named RETAIL. Noticed the chargeback capacity is .89 GB.

<u>Storage System</u>	<u>Mounted Storage</u>	<u>Unmounted Storage</u>	<u>Total Storage</u>
000183500570 (Symm48:38	0.89 GB	0 GB	0.89 GB

Figure 76 Chargeback Capacity

If you were to view the Oracle instance RETAIL in Capacity Manager, you would be shown the local and network capacity, which is a total of 1,042 MB, as shown in the following figure. Of the 1,042 MB, 133 MB is on a local drive. The rest (909 MB) is on a network drive. Convert 909 MB to gigabytes (0.887 GB) and round the output (0.89 GB), the capacity in Capacity Manager matches the number in Chargeback Manager.

RETAIL	Database Instance Name	Total Capacity
	RETAIL	1,042 MB
	Database Files	Total Capacity
	INDX	58 MB
	RBS	520 MB
	RETAILSPACE	5 MB
	SYSTEM	264 MB
	TEMP	72 MB
	TOOLS	12 MB
	USERS	108 MB
	RedoGroup 1	1 MB
	RedoGroup 2	1 MB
	RedoGroup 3	1 MB

Figure 77 Capacity in Capacity Manager

How a Depreciation Method Is Calculated

This section describes the following:

- “Calculating Straight Line Depreciation” on page 477

- ["Calculating Fixed Declining Balance"](#) on page 478
- ["Calculating Double Declining Balance"](#) on page 479

Calculating Straight Line Depreciation

When the management server calculates straight line depreciation, it calculates depreciation based on a fixed rate. These instructions describe how the management server performs the straight line depreciation calculation. An example is provided for each step so that you can try the calculations for yourself.

The following is how the management server calculates straight line depreciation:

1. The management server rolls back the purchase date to the beginning of the purchase month. If the purchase date is later than today (for example, a future purchase), then the purchase date is roll back to today.
Example: Assume the purchase date of an element is January 15, 2003. The management server adjusts the purchased date to January 1, 2003 when calculating months to depreciate.
2. It determines the period ending date. This is equivalent to the last day of the previous full month.
Example: Assume today's date is January 9, 2004. The management server sets the period ending to December 31, 2003.
3. The management server calculates the delta between purchase date and the period ending. This determines how many months worth of depreciation amount the management server need to take into account.
Example: Using the examples from the previous two steps, the delta is 12 months (January 1, 2003 - December 31, 2003).
4. It subtracts the salvage value from the purchase price. This is the depreciable amount.
Example: Assume the purchase price for the element is \$2500 and the Salvage Value is \$100. The depreciable amount is \$2400, which was calculated by subtracting the Salvage Value (\$100) from the purchase price (\$2500).
5. It takes the depreciable amount and divides it by the depreciation period (the number of months it takes for the asset to fully depreciated to either 0 or salvage value). This gives us the depreciation for a single month.
Example: Let's use the depreciable amount (\$2400) calculated in the previous step. Let's assume the depreciation period is 24 months. Divide \$2400 by 24. The result is \$100, which is the one month depreciation.
6. It multiplies the depreciation for a single month by the delta from step 3. This is the total depreciation.
Example: To find the total depreciation, multiply the one month depreciation from the previous step (\$100) by the delta (12 months), which was calculated in Step 3. The result of \$100 x 12 months is \$1,200, which is the total depreciation.
7. To determine the value as of end of last month, the management server simply subtract the total depreciation from the purchase price.
Example: Subtract the total depreciation (\$1200), which was calculated in the last step, from the purchase price (\$2500), which was provided in Step 4. The value as of the end of last month is \$1300.

Calculating Fixed Declining Balance

The Fixed Declining Balance method calculates depreciation based on the value of the asset each month instead of a fixed rate like straight line depreciation.

These instructions describe how the management server performs the fixed declining balance calculation. An example is provided for each step so that you can try the calculations for yourself.

The following is how the management server calculates fixed declining balance:

Example: Assume the purchase date of an element is January 15, 2003. The management server adjusts the purchased date to January 1, 2003 when calculating months to depreciate. The management server rolls back the purchase date to the beginning of the purchase month. If the purchase date is later than today (for example, a future purchase), then the purchase date is roll back to today.

8. It determines the period ending date. This is equivalent to the last day of the previous full month.

Example: Assume today's date is January 9, 2004. The management server sets the period ending to December 31, 2003.

9. The management server calculates the delta between purchase date and the period ending. This determines how many months worth of depreciation amount the management server need to take into account.

Example: Using the examples from the previous two steps, the delta is 12 months (January 1, 2003 - December 31, 2003).

10. The management server takes the user-specified depreciation period and use it as the life of the asset.

Example: Let's assume the depreciation period is 24 months and that it is also the life of the asset.

11. The management server calculates the declining ratio using this formula: $(1.0 / \text{life})$. This determines the rate at which depreciation should occur each month.

Example: Use the example from step 4 (24 months) and use it in the following formula to find the rate of depreciate per month:

$$1.0 / 24$$

The depreciation ratio is 0.042.

12. For each month identified by delta from Step 3, the management server calculates the following:

The example for the following steps can be found at the end of these instructions.

- a. Determine the "would-be" depreciation for the month. This means multiplying the asset value for the month by the declining ratio from step 5.
- b. Subtract the depreciation for the month from the asset value for the month. If the result is less than the salvage value, it means the asset value after depreciation would be less than the salvage. In this case, the management server simply depreciate the asset to the salvage value. Once the management server depreciates an asset down to its salvage value, the depreciation for that asset stops.
- c. If the management server subtracts the depreciation for the month from the asset value and the result is greater than the salvage value, then the management server know it is safe to

depreciate the asset by the depreciation amount calculated in step a. The depreciated asset value for the month would be asset value minus depreciation. The new asset value will be used to compute the depreciation for next month. This process continues until one of the following occurs:

- The management server has depreciated the asset value for the number of months equal to delta.
- The asset value has depreciated down to the salvage cost. If no salvage cost is specified, then the asset value has depreciated down to 0.

Example: For Step 6, let's complete Steps a through c for the first month and then repeat these steps for the second month.

Step 6a - Let's assume the asset value of the element is \$2500. Calculate the "would-be" depreciation of the month by multiplying the asset value by the declining ratio from Step 5 (0.042):

$$\$2500 \times .042 = \$105$$

Step 6b - Assume the salvage value is \$100. Determine if the asset value after depreciation is less than the salvage value by using the following formula.

$$\text{Asset value of the month } (\$2500) - \text{Depreciation for the month } (\$105) = \$2395$$

Since \$2395 (the depreciated asset value) is greater than the salvage value (\$100), the asset value after depreciation is \$2395. Go to Step 6c.

Step 6c - The new asset value (\$2395) is used to calculate the depreciation for the next month. Let's go through the calculations for the next month.

Step 6a - Assume the asset value of the element is \$2395. Calculate the "would-be" depreciation of the month by multiplying the asset value by the declining ratio from Step 5 (0.042):

$$\$2395 \times .042 = \$100.59$$

Step 6b - Assume the salvage value is \$100. Determine if the asset value after depreciation is less than the salvage value by using the following formula:

$$\text{Asset value of the month } (\$2395) - \text{Depreciation for the month } (\$100.59) = \$2294.41$$

Since the \$2294.41 (the depreciated asset value) is greater than the salvage value (\$100), the asset value for the month is \$2294.41. Go to Step 6c. The management server repeats Steps 6a through 6c for 12 months (the delta from Step 3), unless the depreciated asset value reaches the salvage value or 0 if the salvage value is not specified.

Calculating Double Declining Balance

The Double Declining Balance method and the Fixed Declining Balance are very similar. The difference is that instead of using the depreciation ratio determined by $(1.0 / \text{life})$, the management server doubles the ratio to increase the rate of depreciation. This provides for a more realistic depreciation when your asset tends to lose its value in the early part of its life. For instance, a new car's blue book value decreases dramatically once it is sold and driven off the lot of the car dealership.

These instructions describe how the management server performs the double declining balance calculation. An example is provided for each step so that you can try the calculations for yourself:

1. The management server rolls back the purchase date to the beginning of the purchase month. If the purchase date is later than today (for example, a future purchase), then the purchase date is roll back to today.

Example: Assume the purchase date of an element is January 15, 2003. The management server adjusts the purchased date to January 1, 2003 when calculating months to depreciate.

2. It determines the period ending date. This is equivalent to the last day of the previous full month.

Example: Assume today's date is January 9, 2004. The management server sets the period ending to December 31, 2003.

3. The management server calculates the delta between purchase date and the period ending. This determines how many months worth of depreciation amount the management server need to take into account.

Example: Using the examples from the previous two steps, the delta is 12 months (January 1, 2003 - December 31, 2003).

4. The management server takes the user-specified depreciation period and use it as the life of the asset.

Example: Let's assume the depreciation period is 24 months and that it is also the life of the asset.

5. The management server calculates the declining ratio using this formula: $(1.0 / \text{life}) * 2$. This determines the rate at which depreciation should occur each month.

Example: Use the example from step 4 (24 months) and use it in the following formula to find the rate of depreciate per month:

$$(1.0 / 24) * 2$$

The depreciation ratio is 0.084.

6. For each month identified by delta from Step 3, the management server calculates the following:

The example for the following steps can be found at the end of these instructions.

- a. Determine the "would-be" depreciation for the month. This means multiplying the asset value for the month by the declining ratio from step 5.
- b. Subtract the depreciation for the month from the asset value for the month. If the result is less than the salvage value, it means the asset value after depreciation would be less than the salvage. In this case, the management server simply depreciate the asset to the salvage value. Once the management server depreciates an asset down to its salvage value, the depreciation for that asset stops.
- c. If the management server subtracts the depreciation for the month from the asset value and the result is greater than the salvage value, then the management server know it is safe to depreciate the asset by the depreciation amount calculated in step a. The depreciated asset value for the month would be asset value minus depreciation. The new asset value will be used to compute the depreciation for next month. This process continues until one of the following occurs:
 - The management server has depreciated the asset value for the number of months equal to delta.
 - The asset value has depreciated down to the salvage cost. If no salvage cost is specified, then the asset value has depreciated down to 0.

Example: For Step 6, let's complete Steps a through c for the first month and then repeat these steps for the second month.

Step 6a - Let's assume the asset value of the element is \$2500. Calculate the "would-be" depreciation of the month by multiplying the asset value by the declining ratio from Step 5 (0.084):

$$\$2500 \times 0.084 = \$210$$

Step 6b - Assume the salvage value is \$100. Determine if the asset value after depreciation is less than the salvage value by using the following formula.

$$\text{Asset value of the month } (\$2500) - \text{Depreciation for the month } (\$210) = \$2290$$

Since \$2290 (the depreciated asset value) is greater than the salvage value (\$100), the asset value after depreciation is \$2290. Go to Step 6c.

Step 6c - The new asset value (\$2290) is used to calculate the depreciation for the next month. Let's go through the calculations for the next month.

Step 6a - Assume the asset value of the element is \$2290. Calculate the "would-be" depreciation of the month by multiplying the asset value by the declining ratio from Step 5 (0.084):

$$\$2290 \times .084 = \$192.36$$

Step 6b - Assume the salvage value is \$100. Determine if the asset value after depreciation is less than the salvage value by using the following formula:

$$\text{Asset value of the month } (\$2290) - \text{Depreciation for the month } (\$192.36) = \$2097.64$$


Since the \$2097.64 (the depreciated asset value) is greater than the salvage value (\$100), the asset value for the month is \$2097.64. Go to Step 6c. The management server repeats Steps 6a through 6c for 12 months (the delta from Step 3), unless the depreciated asset value reaches the salvage value or 0 if the salvage value is not specified.

Viewing Chargeback

This section describes the following:


- ["Viewing Chargeback by Element"](#) on page 481
- ["Viewing Chargeback by Department"](#) on page 482
- ["Viewing Chargeback by Owner"](#) on page 483

Viewing Chargeback by Element

You can view chargeback for an element by clicking the  icon next to the element listed on the Asset tab. If you see empty values, make sure chargeback has been set up, as described in ["Setting Up Asset-Based Chargeback Manager"](#) on page 468 and ["Setting Up Storage-Based Chargeback Manager"](#) on page 471.

IMPORTANT: The management server displays chargeback information up to the end of the previous month. For example, assume you view chargeback information in the middle of February. The calculations for chargeback would include the month of January, but not February.

To view chargeback by element:


1. Access Chargeback Manager as described in "[Accessing Chargeback Manager](#)" on page 459.
2. Click the  icon next corresponding to the element you want to view chargeback. Asset-based chargeback is displayed.
3. (Applications only) To view storage-based chargeback, click **Storage-based** under the Chargeback node to the left of the Asset-based chargeback information. Chargeback information for the element is displayed. If you see missing values, set up chargeback as described in the topics, "[Setting Up Asset-Based Chargeback Manager](#)" on page 468 and "[Setting Up Storage-Based Chargeback Manager](#)" on page 471.

Viewing Chargeback by Department

You can determine how much a department is being charged for equipment use by viewing chargeback by department. This feature lets you view the monthly costs associated with using hardware and applications. If you see empty values, make sure chargeback has been set up, as described in "[Setting Up Asset-Based Chargeback Manager](#)" on page 468 and "[Setting Up Storage-Based Chargeback Manager](#)" on page 471.

IMPORTANT: The management server displays chargeback information up to the end of the previous month. For example, assume you view chargeback information in the middle of February. The calculations for chargeback would include the month of January, but not February.

To view chargeback by department:

1. Access Chargeback Manager as described in "[Accessing Chargeback Manager](#)" on page 459.
2. Click the **Departments** tab in the right pane.
3. Click the  icon next corresponding to the element you want to view chargeback. You are shown the following information:
 - **Department Name** - Provided when the department was added.
 - **Department Number** - Provided when the department was added.
 - **E-mail** - May be blank if information was not provided.
 - **Phone** - May be blank if information was not provided.
 - **Monthly Infrastructure Cost for Asset** - How much it costs to operate the element on a monthly basis.
 - **Ownership Cost** - How much it costs the department in operating the element

- **Total Asset-based Chargeback** - How much it costs the department in operating the element. This number is based on the following formula:

$$(\text{Monthly Infrastructure cost}) + ((\text{Depreciation}) \times (\text{Ownership \%}))$$
- **Monthly Infrastructure Cost for Storage-based Chargeback** - How much it costs for an application to use a certain amount of storage.
- **Total Storage-based Chargeback** - How much its costs the department for an application to use a certain amount of storage:

Monthly Infrastructure Cost + Ownership Cost
 where Ownership Cost is $(\text{Ownership \%}) \times (\text{Storage Cost})$

This page also displays two tables:

- Table for Asset-Based Chargeback lists the asset, depreciation, ownership percentage, and ownership cost.
- Table for Storage-Based Chargeback lists the application, storage allotted, storage cost, ownership %, and ownership cost. The storage allotted value includes mounted and unmounted storage. Any volumes the application can access are included in the storage calculations.

Viewing Chargeback by Owner

You can view Chargeback for all elements by using the **Ownership** tab. The **Ownership** tab shows the ownership distribution across different departments and helps you with quickly identifying the assets without a department owner.

To view Chargeback by owner:

1. Access the management server.
2. Click **Chargeback** in the left pane.
3. Click the **Ownership** tab in the right pane.
4. Select one of the following from the Chargeback Method drop-down menu:
 - **Asset-based** - Displays chargeback information for assets.
 - **Storage-based** - Displays chargeback information for storage (applications only).

The management server displays asset-based or storage-based chargeback information based on your selection. The management server displays chargeback information from the previous month. For example, assume you view chargeback information in the middle of February. The calculations for chargeback would include the month of January, but not February.

If you see empty values, make sure chargeback has been set up, as described in "[Setting Up Asset-Based Chargeback Manager](#)" on page 468 and "[Setting Up Storage-Based Chargeback Manager](#)" on page 471.

NOTE: You can sort elements according to a column heading. Just click a column heading in the table to sort the data. The arrow next to column heading displays whether the items are being sorted in ascending or descending order. If the arrow is pointing up, items are sorted in ascending and alphabetical order. If the arrow is pointing down, items are sorted in descending and reserved alphabetical order.

The following information is displayed:

- **Department Owner** - The department that owns the element. This information was provided when the department was added.
- **Department Number** - The number of the department that owns the element. This information was provided when the apartment was added.
- **Asset Name** - The name of the element.
- **Vendor** - The company that supplied the element.
- **Serial Number** - The serial number of the element.
- **Ownership Percentage** - The percentage the department owns of the element.
- **Storage (GB)**** - The amount in gigabytes the application uses. This value includes mounted and unmounted storage. Any volumes the application can access are included in the storage calculations.
- **Storage Cost**** - (Storage-based only) How much it costs to run the storage that the application uses.
- **Ownership Cost** - How much it costs the department to use the asset.

**Applies to only storage-based chargeback.

Chargeback Reports

This section describes the following:

- ["Viewing Chargeback Reports"](#) on page 484
- ["E-mailing a Chargeback Report"](#) on page 485
- ["Managing E-mail Schedules for Chargeback Reports"](#) on page 486

Viewing Chargeback Reports

You can access Chargeback reports one of two ways:

- From Chargeback Manager
- From Reporting

Keep in mind the following:

- To populate chargeback reports, enter information for chargeback, as described in the topic, ["Setting Up Chargeback Manager"](#) on page 458.
- If you want to view the latest information in a report, click the **Refresh View** button in the upper-right corner of Reporter. The management server refreshing the page displaying the report.

To view Chargeback reports:

1. Do one of the following:
 - To access the reports through Chargeback, click the **Reports** tab in Chargeback Manager and expand the **Chargeback Reports** node in the tree in the middle pane.
 - To access the reports through Reporter, click the **Reporter** button. Then, expand the **Reporter > Chargeback** nodes in the tree in the middle pane.
2. Click one of the default reports:
 - **Storage System-Based Assets** - Displays asset information from storage arrays: host name, department, HBA port, HBA port WWN, storage volume, volume size, and cost.
 - **Asset Based** - Displays the following asset-based chargeback for each department owning elements: department, asset name, ownership ratio, and chargeback amount. Total asset-based cost per month is also displayed.
 - **Storage Based** - Displays storage-based chargeback for each department owning applications: department, application, ownership ratio, total capacity, and capacity amount. Total storage-based cost per month is also displayed.
 - **Storage System Based by Tier** - Displays storage-based chargeback by tier. Only storage systems are assigned storage tiers.

The report is displayed as a Web page in the right pane. See the topic, "[Chargeback Reports](#)" on page 484 for more information about each of the reports.
3. Select one of the following options for a different output:
 - **PDF**
 - **Excel**
 - **XML**
4. Click the **View Report in New Window** link at the top of the page to view the report in a separate window.

E-mailing a Chargeback Report

You can e-mail a chargeback report in PDF, XML or Microsoft Excel format. If you want to e-mail reports by on a regular basis, set up an e-mail schedule for the report, as described in "[Adding an E-mail Schedule for a Chargeback Report](#)" on page 486.

IMPORTANT: Before you can e-mail a report, you must set up e-mail notification, as described in the topic, "[Setting Up E-mail Notification](#)" on page 99.

To e-mail a report:

1. Do one of the following:
 - To access the reports through Reporter, click the **Reporter** button. Then, expand the **Reporter > Chargeback** nodes in the tree in the middle pane.
 - To access the reports through Chargeback, click the **Reports** tab in Chargeback Manager and expand the **Chargeback Reports** node in the tree in the middle pane.

2. Expand the tree in the middle pane, and double-click the report for which you want to send by e-mail.
3. When the report is displayed in the right pane, click the **E-mail Schedule** tab in the right pane.
4. Click the **E-mail Report** button.
5. In the **To** field, type the recipient's e-mail address.
The software verifies the address entered has a correct form. To send multiple addresses, separate each address with a semicolon (;), for example:
`john.example@appiq.com;jerry.example@appiq.com`
6. From the **Format** drop-down menu, select one of the following formats:
 - **PDF** - Requires the use of Adobe Acrobat, which can be downloaded for free from <http://www.adobe.com>.
 - **Excel** - Requires the use of Microsoft Excel.
 - **XML** - Requires the user has an understanding of XML.
7. (Optional) Modify the subject and message.
8. Click the **OK** button.
The report is sent.

Managing E-mail Schedules for Chargeback Reports

This section describes the following:

- ["Adding an E-mail Schedule for a Chargeback Report"](#) on page 486
- ["Editing an E-mail Schedule for a Chargeback Report"](#) on page 488
- ["Deleting E-mail Schedules for a Chargeback Report"](#) on page 489
- ["Viewing E-mail Schedules for a Chargeback Report"](#) on page 490
- ["Viewing the History of an E-mail Chargeback Schedule"](#) on page 491

Adding an E-mail Schedule for a Chargeback Report

You can add an e-mail schedule so that a user receives an attached report on a regular basis. The report can be in the form a PDF, XML or Microsoft Excel document.

Keep in mind the following:

- Before you can add an e-mail schedule, you must set up e-mail notification, as described in the topic, ["Setting Up E-mail Notification"](#) on page 99.
- The management server service must be running for users to receive e-mail notification.
- Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager in the upper-right corner of the screen. Then, click the **Scheduled Deliveries** tab at the top of the screen.

To add an e-mail schedule:

1. Access the management server.

2. Do one of the following:
 - To access the reports through Chargeback, click the **Reports** tab in Chargeback Manager and expand the **Chargeback Reports** node in the tree in the middle pane.
 - To access the reports through Reporter, click the **Reporter** button. Then, expand the **Reporter > Chargeback** nodes in the tree in the middle pane.
3. Click the report for which you want to set up an e-mail schedule.

The report is displayed as a Web page in the right pane. See the topic, "[Chargeback Reports](#)" on page 484 for more information about each of the reports.
4. When the report is displayed in the right pane, click the **E-mail Schedule** tab in the right pane.
5. Click the **Add E-mail Schedule** button.
6. In the **To** field, type the recipient's e-mail address.

The software verifies the address entered has a correct form. To send multiple addresses, separate each address with a semicolon (;), for example:

john.example@appiq.com;jerry.example@appiq.com
7. In the **Subject** field, type a subject for the e-mail messages you plan to send.

NOTE: Provide the name of the report in the subject field so users can distinguish the message from others.

8. In the **Message** field, type a message describing the report.

If you are e-mailing reports in bulk, you might want to let users know the e-mail is being sent by an automated process. You might also want to provide an e-mail address for users to provide feedback, for example:

This e-mail and its attached report are generated automatically. If you would like to change how often the report is sent to you or you want to be taken off the list, please contact username@companyname.com.
9. From the **Format** drop-down menu, select one of the following formats:
 - **PDF** - Requires the use of Adobe Acrobat, which can be downloaded for free from <http://www.adobe.com>.
 - **Excel** - Requires the use of Microsoft Excel.
 - **XML** - Requires the user has an understanding of XML.
10. In the **Time to Run** field, type the time you want to send the report. This time must be entered in the 24-hour format. For example, if you want a report sent at 2:15 p.m., you would type 14:15 in **Time to Run** field.
11. Select one of the following options to determine how frequently you want to send the report.
 - **Daily** - If you selected daily, select how frequently you want the management server to send the report.
 - **Everyday** - The report is sent everyday.
 - **Weekday** - The report is sent only Monday through Friday.
 - **Everyday for a specified number of days** - Fill in the number of days you want the report to be sent daily. After the specified number of days, the report is no longer sent.

For example, you could use this feature to send reports to someone's replacement while the person is away on vacation.

- **Weekly** - If you selected weekly, use the **Frequency** drop-down menu to select the day of the week on which you want the report sent.
- **Monthly** - If you selected monthly, select the time during the month you want the report sent.
 - To send the report on the first or last day of the month, select the first option. Then, select **First** or **Last** from the drop-down menu.
 - To send the report on a specified day during the month, select the second option. Then, type the day on which you want the report sent. Keep in mind that the number of days in a month vary. So if you type 30 in this field, users will not receive a report in February. Also, if you type 29 in this field, users do not receive the report in February during non-leap years.

12. Click the **OK** button.


The schedule is created.

Editing an E-mail Schedule for a Chargeback Report

IMPORTANT: Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager in the upper-right corner of the screen. Then, click the **Scheduled Deliveries** tab at the top of the screen.

To edit an e-mail schedule for a report:

1. Access the management server.
2. Do one of the following:
 - To access the reports through Chargeback, click the **Reports** tab in Chargeback Manager and expand the **Chargeback Reports** node in the tree in the middle pane.
 - To access the reports through Reporter, click the **Reporter** button. Then, expand the **Reporter > Chargeback** nodes in the tree in the middle pane.
3. Click the report for which you want to set up an e-mail schedule.

The report is displayed as a Web page in the right pane. See the topic, "[Chargeback Reports](#)" on page 484," for more information about each of the reports.
4. When the report is displayed in the right pane, click the **E-mail Schedule** tab in the right pane.
5. Under the Edit column, click the **Edit** () button.
6. In the **To** field, change the recipient's e-mail address.

The software verifies the address entered has a correct form. To send multiple addresses, separate each address with a semicolon (;), for example:

john.example@appiq.com;jerry.example@appiq.com
7. In the **Subject** field, change the subject of the e-mail.
8. In the **Message** field, change a message describing the report.

If you are e-mailing reports in bulk, you might want to let users know the e-mail is being sent by an automated process. You might also want to provide an e-mail address for users to provide feedback, for example:

This e-mail and its attached report are generated automatically. If you would like to change how often the report is sent to you or you want to be taken off the list, please contact username@companyname.com.

9. From the **Format** drop-down menu, select one of the following formats:

- **PDF** - Requires the use of Adobe Acrobat, which can be downloaded for free from <http://www.adobe.com>.
- **Excel** - Requires the use of Microsoft Excel.
- **XML** - Requires the user has an understanding of XML.

10. In the **Time to Run** field, type the time you want to send the report. This time must be entered in the 24-hour format. For example, if you want a report sent at 2:15 p.m., you would type 14:15 in **Time to Run** field.

11. Select one of the following options to determine how frequently you want to send the report.


- **Daily** - If you selected daily, select how frequently you want the management server to send the report.
 - **Everyday** - The report is sent everyday.
 - **Weekday** - The report is sent only Monday through Friday.
 - **Everyday for a specified number of days** - Fill in the number of days you want the report to be sent daily. After the specified number of days, the report is no longer sent. For example, you could use this feature to send reports to someone's replacement while the person is away on vacation.
- **Weekly** - If you selected weekly, use the **Frequency** drop-down menu to select the day of the week on which you want the report sent.
- **Monthly** - If you selected monthly, select the time during the month you want the report sent.
 - To send the report on the first or last day of the month, select the first option. Then, select **First** or **Last** from the drop-down menu.
 - To send the report on a specified day during the month, select the second option. Then, type the day on which you want the report sent. Keep in mind that the number of days in a month vary. So if you type 30 in this field, users will not receive a report in February. Also, if you type 29 in this field, users do not receive the report in February during non-leap years.

12. Click the **OK** button.

Deleting E-mail Schedules for a Chargeback Report

IMPORTANT: Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, click **Reports > Storage Essentials > Data Collection > System Data** in HP Systems Insight Manager in the upper-right corner of the screen. Then, click the **Scheduled Deliveries** tab at the top of the screen.

To delete an e-mail schedule:

1. Access the management server.
2. Do one of the following:
 - To access the reports through Chargeback, click the **Reports** tab in Chargeback Manager and expand the **Chargeback Reports** node in the tree in the middle pane.
 - To access the reports through Reporter, click the **Reporter** button. Then, expand the **Reporter** > **Chargeback** nodes in the tree in the middle pane.
3. Click the report for which you want to set up an e-mail schedule.
The report is displayed as a Web page in the right pane. See the topic, "[Chargeback Reports](#)" on page 484 for more information about each of the reports.
4. Click the  button corresponding to the e-mail schedule you want to remove.

Viewing E-mail Schedules for a Chargeback Report

IMPORTANT: Only the e-mail schedules created by the current user are listed. To view the e-mail schedules for all reports, click **Reports** > **Storage Essentials** > **Data Collection** > **System Data** in HP Systems Insight Manager in the upper-right corner of the screen. Then, click the **Scheduled Deliveries** tab at the top of the screen.




To view the E-mail schedules assigned to a report.

1. Access the management server.
2. Do one of the following:
 - To access the reports through Chargeback, click the **Reports** tab in Chargeback Manager and expand the **Chargeback Reports** node in the tree in the middle pane.
 - To access the reports through Reporter, click the **Reporter** button. Then, expand the **Reporter** > **Chargeback** nodes in the tree in the middle pane.
3. Click the report for which you want to set up an e-mail schedule.
4. When the report is displayed in the right pane, click the **E-mail Schedule** tab in the right pane.
Information about the e-mail schedules for that report are displayed.


Table 120 Viewing E-mail Schedules for a Chargeback Report

Column Name	Description
Recipient	The person who receives the report.
Subject	The subject of the e-mail, brief summary of what it is about.
Format	The format of the report sent: <ul style="list-style-type: none"> • PDF • Microsoft EXCEL • XML


Table 120 Viewing E-mail Schedules for a Chargeback Report (continued)

Column Name	Description
Last Delivered	The time the last report was sent to the recipient.
History	Click the View button to display the times and dates when the report was sent. You can also delete a historical entry by clicking the  button for the corresponding entry.
Edit	Click the  button to edit a schedule of the report. See the topic, " Adding an E-mail Schedule for a Chargeback Report " on page 486 for information about the options displayed in this window.
Delete	Click the  button to remove the corresponding schedule.

Viewing the History of an E-mail Chargeback Schedule

You can display the times and dates when the report was sent. You can also delete a historical entry by clicking the  button for the corresponding entry.

To view the history of an e-mail schedule:

1. Access the management server.
2. Do one of the following:
 - To access the reports through Chargeback, click the **Reports** tab in Chargeback Manager and expand the **Chargeback Reports** node in the tree in the middle pane.
 - To access the reports through Reporter, click the **Reporter** button. Then, expand the **Reporter** > **Chargeback** nodes in the tree in the middle pane.
3. Click the report for which you want to set up an e-mail schedule.
4. When the report is displayed in the right pane, click the **E-mail Schedule** tab in the right pane.
5. Under the History column, click the **View** button.
You are shown when the report was sent.
6. To remove a historical entry, click the  button.

Filtering Assets

This section describes the following:

- "[About Filtering Assets](#)" on page 492
- "[Selecting an Element Type for Chargeback](#)" on page 492
- "[Customizing the Element Type Filter for Chargeback](#)" on page 493
- "[Filtering Assets by Status](#)" on page 493

- ["Customizing the Asset Status Filter for Chargeback"](#) on page 493
- ["Hiding Filters in Chargeback Manager"](#) on page 494

About Filtering Assets

The management server provides several types of filters to specify which assets you want Chargeback to display.

You can use all the filters at once or you can use just one of them. You can filter assets by:

- **Status**
- **Element type**

For example, assume you need a host to install an application, but you are not too sure which hosts are in use. You could set the filters so that only hosts with a status of "in use" are displayed. You could then click the element to find contact information for the owner.

You could also use the filters to find out which elements are missing or repaired. For example, assume you want to find out which hosts are missing or repaired.

You could do the following:

- Set the filter to display only hosts by selecting the **Host** option from the Show Element Type combo box.
- Click the **Custom** button next to the Show Status combo box. Verify that only Missing and Repaired are selected. Click **OK**.

Once you set all of your filters, you would click the **Apply Filters** button.

Selecting an Element Type for Chargeback

You can filter by element type, so only certain types of elements are displayed. For example, you can specify that only hosts are displayed.

To filter by element type, select an option from the **Show Element Type** drop-down menu in Chargeback. When you are asked if you want to apply your changes, click **Yes**. If you want to apply your changes at a later time you can click **No** and then, click the **Apply Filters** button when you are ready for your changes to take effect. Chargeback displays only the elements you specified in your filter.

Table 121 Element Types

Element Type	Description
Applications	Displays only applications, such as Microsoft Exchange and Oracle.
Host	Displays only hosts.
Switch	Displays only switches.
Storage System	Displays only storage systems.

Table 121 Element Types (continued)

Element Type	Description
All	Lists all elements.

Customizing the Element Type Filter for Chargeback

You can customize the element type for your filter by clicking the **Customize** button next to the **Show Element Type** drop-down menu in Chargeback.

For example, you can specify you want only hosts and switches displayed in Chargeback.

To select more than one element for filtering:

1. Click the **Custom** button next to the **Show Element Type** drop-down menu in Chargeback.
2. Select the element types you want displayed in Chargeback.
3. Click **OK**.
4. When you are asked if you want to apply your changes, click **Yes**. If you want them to take effect immediately. If not, click **No**. Then, click the **Apply Filters** button when you want your changes to take effect.

Chargeback displays the elements that meet the criteria selected in the filter.

Filtering Assets by Status

You can filter an asset by status, so only certain assets of a specified status are displayed. For example, you can specify that only assets in use are displayed.

To filter by asset status:

1. Select the following options from the **Show Element Type** drop-down menu in Chargeback.
 - **All (Default)** - All assets are displayed.
 - **New** - Only assets with the status of New are displayed.
 - **Missing** - Only assets with the status of Missing are displayed.
 - **In Use** - Only assets with the status of In Use are displayed.
2. When you are asked if you want to apply your changes, click **Yes**. If you want to apply your changes at a later time you can click **No** and then, click the **Apply Filters** button when you are ready for your changes to take effect. Chargeback displays only the assets you specified in your filter.

Customizing the Asset Status Filter for Chargeback

You can filter multiple assets by clicking the **Customize** button next to the **Show Status** drop-down menu in Chargeback.

For example, you can specify you want only assets that are missing displayed in Chargeback.

To select more than one assets for filtering:

1. Click the **Custom** button next to the **Show Status** drop-down menu in Chargeback Manager.
2. Select the statuses you want displayed in Chargeback Manager.

3. Click **OK**.
4. When you are asked if you want to apply your changes, click **Yes**. if you want them to take effect immediately. If not, click **No**. Then, click the **Apply Filters** button when you want your changes to take effect.

Chargeback Manager displays the assets that meet the criteria selected in the filter.

Hiding Filters in Chargeback Manager

Hide the filters for additional screen space. When you hide the filters, the following features are hidden:

- Show Element Type
- Show Status

To hide the filters, click the **Hide Filters** link in the upper-left corner of Chargeback Manager.

To display the filters, click the **Show Filters** link in the upper-left corner of Chargeback Manager.

17 Business Tools

IMPORTANT: Depending on your license, Business Tools may not be available. See the “List of Features” to determine if you have access to Business Tools. The “List of Features” is accessible from the Documentation Center (**Help > Documentation Center** in Storage Essentials).

This chapter describes the following:

- “[About the Business Tools](#)” on page 495
- “[Using the HBA Replacement Automator](#)” on page 496
- “[Using the Brocade Firmware Update Utility](#)” on page 497
- “[Setting up Risk Analysis](#)” on page 499
- “[Global Change Management](#)” on page 499

About the Business Tools

Business Tools provides the following functionality to help you manage the business aspect of your network.

- **Advisors** - Provides detailed information for you to make informed decisions about your network, such as non-compliant HBA firmware. You may have access to the following advisors depending on your release:
 - **Reachable Storage** - Provides information about the storage accessible from a selected host, such as the following:
 - Free volumes on current storage systems
 - LUNs mapped to host, but not mounted with file systems
 - Free volumes on other storage systems in host fabrics
 - Free volumes on all other storage systems
 - **HBA Risk Analysis** - It examines whether the HBAs are at risk. Configure the `hba_risk_analysis.conf` file before you run this tool. See “[Setting up Risk Analysis](#)” on page 499 for more information.
 - **Switch Risk Analysis** - It examines whether the switches are at risk. Configure the `switch_risk_analysis.conf` file before you run this tool. See “[Setting up Risk Analysis](#)” on page 499 for more information.
 - **Global Change Management** - It lets you record the properties and connections of managed elements. See “[Global Change Management](#)” on page 499.
- **Automators** - Lets you automate particular tasks. You may have access to the following automators depending on your release:
 - **Brocade Firmware Update Utility** - Lets you update Brocade firmware remotely.
 - **HBA Replacement Automator** - Makes the management server aware of a replaced HBA card so that the latest information for zoning and LUN masking is still available. See “[Using the HBA Replacement Automator](#)” on page 496 before you use this tool.

Keep in mind the following:

- After you use an automator, refresh the element or perform Discovery Data Collection so the management server can obtain new information and update the user interface accordingly. Only Discovery Data Collection removes elements that are no longer there from the user interface. For example, removed ports could appear in Port Details but not in the topology for the host if you refresh the element, instead of Discovery Data Collection.
- All discovered elements are accessible in Business Tools, regardless of a user's restrictions. For example, assume your account belongs to an organization that has only hosts as members. If you run Switch Risk Analysis, the management server still provides information about whether the switches are a risk in your environment.

Using the HBA Replacement Automator

The HBA Replacement Automator makes the management server aware of a replaced HBA card so that the latest information for zoning and LUN masking is still available. You must complete all steps to prevent LUN masking from being lost after running the HBA Replacement Automator.

IMPORTANT: After installing the new HBA, you must run Discovery Data Collection before you run the HBA Replacement Automator script; otherwise, you cannot change zone and zone aliases to accommodate the new Worldwide Name (WWN).

You have two ways to replace the HBA. You can install the new HBA with the old HBA or you can install the new HBA by itself.

Installing New HBA with Old HBA

To install the new HBA with the old HBA:

1. Install the new HBA with the old HBA.
2. Run Discovery Data Collection.
3. Run the HBA Replacement Automator (**Tools > Storage Essentials > Business Tools**).
4. When you run the Automator, select the WWNs of the new and old HBA.

Installing New HBA by Itself

To install the new HBA by itself:

1. Write down the WWN of the HBA you want to replace. You can find the WWN name of the HBA by doing the following:
 - a. Double-click the host of the HBA in System Manager.
 - b. In the Navigation pane, click the **Host Bus Adapters** button at the top of the pane.
The host bus adapters for the host are listed and their WWNs. Write down the WWN.
2. Install the new HBA by itself.
3. Run Discovery Data Collection (**Discovery > Details**).
4. Run the HBA Replacement Automator (**Business Tools > HBA Replacement Automator**).
5. When you run the Automator, enter the WWN of the old HBA.

Using the Brocade Firmware Update Utility

You can update the firmware of a Brocade switch remotely by using the Brocade Firmware Update Utility on the management server. The management server must obtain some information about the switch and the location of the firmware before it can upgrade the firmware on the switch. This information can be entered in a configuration file or it can be entered when you run the Brocade Firmware Update Utility.

The directory containing the new firmware must be:

- in a shared network folder
- on an server with FTP capabilities

Updating the Firmware Without Modifying the Configuration File

To update the firmware without modifying the configuration file:

1. Click **Tools > Storage Essentials > Business Tools**.
2. Click **Brocade Firmware Update Utility** in the right pane under Automators.
3. Select the switch you want to upgrade, and then click **OK**.

If you have not provided in the configuration file the user name and password of the switch, you are shown the following message.

One or both of the following keys are missing or contain invalid values:

```
'switch_user'  
'switch_pass'
```

4. In the **Commands** field, type the user name of the switch. Press ENTER.
5. In the **Commands** field, enter the password of the switch.

If you have not provided in the configuration file the user name and password of the server that has the firmware, you are shown the following message.

One or both of the following keys are missing or contain invalid values:

```
'ftp_user'  
'ftp_pass'
```

6. In the **Commands** field, enter the user name of the server that has the firmware.
7. In the **Commands** field, enter the password of the server that has the firmware.

If you have not provided in the configuration file the IP address of the server that has the firmware, you are shown the following message:

```
""is not a valid value for key 'ftp_server'
```

8. In the **Commands** field, enter the IP address of the server that has the firmware.

If you have not provided the path to the firmware, you are shown the following message:

```
"" is not a valid value for key 'firmware'.
```

9. Enter the complete path to the firmware. Include the file name in the path, but not the drive letter. Use slashes (/) even if the firmware is on a Microsoft Windows computer, for example:
`/myshare/firmware_directory/firmware_filename.`

The directory containing the new firmware must be:

- in a shared network folder

- on an server with FTP capabilities
10. When you are asked if you want to upgrade, enter **Yes**.
 11. After the firmware is loaded, reboot the switch.

Updating the Firmware With Modifying the Configuration File

You can enter the information for the firmware upgrade in the `update_Brocade.conf` file located in `%JBOS4_DIST%\server\appiq\remotescripts\automators` where `[Install_Dir]` is the directory into which you installed the management server.

1. Open the `update_Brocade.conf` file in a text editor, such as Notepad.
2. Acceptable fields are the following:
 - `switch_user` - username for switch
 - `switch_pass` - password for switch username
 - `ftp_server` - IP address of FTP server with firmware file
 - `ftp_user` - username for FTP server
 - `ftp_pass` - password for FTP username
 - `firmware` - absolute path to firmware file on FTP server
3. Assigned the following values. Verify the number sign (#) is removed from in front of the field. The program will prompt for these fields if they are not specified. All fields must be tab delimited.

```
switch_user admin
switch_pass switchpassword
ftp_server 192.168.1.1
ftp_user root
ftp_pass password
firmware /myshare/firmware_directory/firmware_file
where
```

- `admin` is the user name for the switch
 - `switchpassword` is the password for the switch
 - `192.168.1.1` the IP address of the server that has the new firmware
 - `root` is the user name of the server that has the new firmware
 - `password` is the password of the server that has the new firmware
 - `/myshare/firmware_directory/firmware_file` is the full path to the firmware file.
4. Save the `update_Brocade.conf` file.
 5. Click **Tools > Storage Essentials > Business Tools**.
 6. Select the switch you want to upgrade, and then click **OK**.
 7. When you are asked if you want to upgrade the firmware, enter **Yes**.
 8. After the firmware is updated, reboot the switch.

Setting up Risk Analysis

The risk analysis tools flag which HBAs and/or switches are a risk in your environment. The HBA Risk Analysis and Switch Risk Analysis determine which elements are at risk by checking them against predefined profiles you created. Before you run either of these tools, create profiles for your environment, as described in the following steps.

1. Go to the `%JBSS4_DIST%\server\appiq\remotescripts\advisors` directory and open one of the following files in a text editor, such as Notepad:
 - **HBA Risk Analysis** - `hba_risk_analysis.conf`
 - **Switch Risk Analysis** - `switch_risk_analysis.conf`

where `[Install_Dir]` is the directory into which you installed the management server.

2. In the configuration file, do one of the following:
 - If you are editing `hba_risk_analysis.conf`, define a profile for each type of HBA in your environment.
 - If you are editing `switch_risk_analysis.conf`, define a profile for each type of switch in your environment.

The profile name can be any name you want, and the number and type of fields in the profile is defined by you. For example, assume you want to check that all HBAs in Solaris hosts are a certain model (LP-9000). The profile would be in `hba_risk_analysis.conf` and it would resemble the following:

```
any profile name
the 'OS' field = Solaris
the 'Model' field = LP-9000
```

Let's expand that analysis to verifying that those HBA's also have a certain driver version and a certain firmware level, as shown in the following example profile:

```
any profile name
the 'OS' field = Solaris
the 'Model' field = LP-9000
the 'DriverVersion' field = whatever
the 'FirmwareVersion' field = whatever
```

IMPORTANT: The use of 'special' characters in the field values are interpreted as regular expressions (or wildcards) by the search. For example, "Model LP900[02]" will match the values "LP-9000" or "LP-9002". "FirmwareVersion 3.90A[0-9]" will match any firmware version 3.90A0 through 3.90A9.

3. Once you have your profiles configured, run HBA Risk Analysis or Switch Risk Analysis. The tool compares every HBA or switch against the profiles and flag any HBA or switch that does not match at least one profile.

Global Change Management

Global Change Management lets you save the current configuration and/or compare changes with a previous configuration.

Accessing Global Change Management

1. Access Business Tools.
2. Click **Global Change Management Tools**.

Saving the Current Configuration

To save the current configuration, use option 0. The current configuration is saved to a DAT file in the following directory. This file can be opened by using a text editor, such as Notepad:

```
%JBoss4_DIST%\server\appiq\remotescripts\advisors\saved-configurations
```

The DAT file contains the data for the elements discovered in the configuration. The following is an example of a portion of a DAT file for a saved configuration:

```
$VAR1 = 'conf3';
$VAR2 = 1117783473;
__DATA__
$VAR1 = {
  '1002' => {
    'PROPERTIES' => {
      'ID' => '1002',
      'HostType' => 'Default',
      'DnsName' => 'QA67',
      'SupportFlags' => '7',
      ....
    }
  }
}
```

where

- \$VAR1 is the configuration name, which is conf3 in this case.
- 1002 is the element ID.

Comparing a Previous Configuration by Using Global Change Management

NOTE: Global Change Management requires you to provide the name of the current configuration the first time you run the tool. Global Change Management then assigns the saved configuration name to a number.

To compare a previous configuration, enter the number corresponding with the previous saved configuration. While the script is determining the changes, it lists the elements it is analyzing. Once it is done, it lists the changes under the heading “CHANGED PROPERTIES” on the screen.

18 Troubleshooting

This chapter describes the following:

- [“Data is late or an error occurred” Message](#) on page 501
- [“appiq.log Filled with Connection Exceptions”](#) on page 501
- [“Receiving “HTTP ERROR: 503” When Accessing the Management Server”](#) on page 502
- [“Permanently Changing the Port a CIM Extension Uses \(UNIX Only\)”](#) on page 503
- [“Configuring UNIX CIM Extensions to Run Behind Firewalls”](#) on page 504
- [“Volume Names from Ambiguous Automounts Are Not Displayed”](#) on page 508
- [“Installing the Software Security Certificate”](#) on page 509
- [“Troubleshooting Discovery and Discovery Data Collection”](#) on page 511
- [“Troubleshooting the Java Plug-in”](#) on page 521
- [“Troubleshooting Provisioning”](#) on page 522
- [“Troubleshooting Hardware”](#) on page 523

“Data is late or an error occurred” Message

If you see the message “Data is late or an error occurred” when you tried to obtain information from a UNIX host, verify you were logged in as root when you started the CIM Extension (`./start`). You must be logged in as root if you want to use the `./start` command, even if you are using the

`./start -users username` command, where `username` is a valid UNIX account.

The CIM Extension only provides the information within the privileges of the user account that started the CIM Extension. This is why you must use root to start the CIM Extension. Only root has enough privileges to provide the information the management server needs.

If you continue to see the message, contact customer support.

appiq.log Filled with Connection Exceptions

When an Oracle REDO log becomes corrupt, the management server is unable to connect to the database. Whenever the management server is unable to connect to the Oracle database, it writes to the `appiq.log` file. Many exceptions may cause the Application Log on Windows to become full.

To fix the problem, stop the management server and Oracle. Then, remove the corrupted REDO log, as described in the following steps:

1. Stop the AppStorManager service, which is the service the management server uses.

NOTE: While the service is stopped, the management server cannot monitor elements and users cannot access the management server.

2. To find the corrupt log file, look in the alert_APPIQ.log file, which can be found in \oracle\admin\APPIQ\bdump, where ORACLE_BASE is c:\oracle
You can verify if the REDO log listed in the alert_APPIQ.log file is corrupt by looking for a “redo block corruption” error in the REDO log.
3. On the management server, enter the following at the command prompt:
`Sqlplus /nolog`
4. Enter the following:
`Sql> connect sys/change_on_install as sysdba`
5. Enter the following:
`Sql> startup mount;`
6. Enter the following:
`Sql> ALTER DATABASE CLEAR UNARCHIVED LOGFILE
'C:\ORACLE\ORADATA\APPIQ\REDO02.LOG';`
where C:\ORACLE\ORADATA\APPIQ\REDO02.LOG is the corrupted log file and its path.
7. Enter the following:
`Sql> alter database open`
8. Enter the following:
`Sql> shutdown immediate;`
9. Enter the following:
`Sql> startup`

Receiving “HTTP ERROR: 503” When Accessing the Management Server

If you receive a message resembling the following when you try to access the management server, make sure your database for the management server is running. If it is not, start the database.

```
Receiving HTTP ERROR: 503 javax.ejb.EJBException: null;
```

Refer to the following sections for more information about how to start database for the management server.

Access the Services window to make sure the OracleOraHome92TNSListener service has started and is set to automatic. Refer to the Windows documentation for information on how to access the Services window.

If the OracleOraHome92TNSListener service has not started but the AppStorManager service has started, start the OracleOraHome92TNSListener service and then restart AppStorManager.

IMPORTANT: If you are starting the services manually, start the Oracle service before the service for the management server.

Errors in the Logs

If you access the logs, you are shown messages resembling the following. The complete text has been shortened as a result of space constraints:

```
Aug 04 2004 11:59:07] INFO
[com.appiq.service.policyManager.policyService.PolicyService] Creating
[Aug 04 2004 11:59:07] INFO
[com.appiq.service.policyManager.policyService.PolicyService] Created
[Aug 04 2004 11:59:07] INFO
[com.appiq.service.policyManager.policyService.PolicyService] Starting
[Aug 04 2004 11:59:07] INFO
[com.appiq.service.policyManager.policyService.PolicyService] Starting
Policy Factory
[Aug 04 2004 11:59:11] ERROR [com.appiq.security.DatabaseSecurityManager]
DatabaseSecurityManager Error:
org.jboss.util.NestedSQLException: Could not create connection; - nested
throwable: (java.sql.SQLException: ORA-01033: ORACLE initialization or
shutdown in progress
); - nested throwable: (org.jboss.resource.ResourceException: Could not
create connection; - nested throwable: (java.sql.SQLException: ORA-01033:
ORACLE initialization or shutdown in progress
))
```

Permanently Changing the Port a CIM Extension Uses (UNIX Only)

CIM Extensions on UNIX use port 4673 by default. You can start a CIM Extension on another port by entering `./start -port 1234`, where 1234 is the new port. With this method, you must always remember to provide the nondefault port when starting the CIM Extension.

You can configure a CIM Extension to remember the nondefault port, so you only need to enter `./start` to start the CIM Extension:

1. Go to the `/opt/APPQcime/tools` directory.
2. In this directory, create a file called `cxws.host.parameters`.
3. Open the newly created file in a text editor, and provide the following line:

```
-credentials username:password
-port 1234
```

IMPORTANT: The values for `-credentials` and `-port` must be on separate lines, as shown in the example.

where

- `username` is the user that is used to discover the CIM Extension. You will need to provide this user name and its password when you discover the host.

- password is the password of username.
 - 1234 the new port for the CIM Extension
4. Save the file.
 5. Restart the CIM Extension for your changes to take effect.

NOTE: The CIM Extension looks for parameters in the `cxws.host.parameters` file whenever it starts, such as when it is started manually or when the host is rebooted.

6. The management server assumes the CIM Extension is running on port 4673. If you change the port number, you must make the management server aware of the new port number. In the IP Address/DNS Name field, type a colon and then the port number after the IP address or DNS name, as shown in the following example:

```
192.168.1.2:1234
```

where

- 192.168.1.2 is the IP address of the host
- 1234 is the new port number

If you have already added the host to the discovery list on the management server, you must remove it and then re-add it. You cannot have more than one listing of the host with different ports.

Configuring UNIX CIM Extensions to Run Behind Firewalls

In some instances you will need to discover a host behind a firewall. Use the following table as a guideline. Assume the management server wants to discover HostA, which has three network interface cards on three separate networks with three separate IPs: 10.250.250.10, 172.31.250.10, and 192.168.250.10. In the following table different configurations are presented:

- The “Manual Start Parameters for CIM Extensions” column provides what you would enter to start the CIM Extension manually on the host. See the Installation Guide for more information on how to start a CIM Extension manually.
- The “If Mentioned in `cxws.host.parameters`” column provides information on how you would modify the `cxws.host.parameters` file. See [“Permanently Changing the Port a CIM Extension Uses \(UNIX Only\)”](#) on page 503.
- The “Step 1 Discovery and RMI Registry Port” column - Provides information about what IP addresses are required for the discovery list. The RMI Registry port is the port the CIM Extension uses. Keep in mind that when a port other than 4673 is used for the CIM Extension, the port must be included in the discovery IP. For example, 192.168.1.1:1234, where 192.168.1.1 is the IP for the host and 1234 is the port the CIM Extension uses.

Table 122 Troubleshooting Firewalls

Configur- ation	Manual Start Parameters for CIM Extension	If Mentioned in cxws.host.parameters	Step 1 Discovery and RMI Registry Port
Firewall port 4673 opened between host and management server	start		10.250.250.10 OR 172.31.250.10 OR 192.168.250.10 Communication Port: 4673
Firewall port 1234 opened between host and management server, but specific port	start -port 1234	-port 1234	10.250.250.10:1234 OR 172.31.250.10:1234 OR 192.168.250.10:1234 Communication Port: 1234
Firewall port 4673 opened between host and management server on the 172.31.250.x subnet	start -on 172.31.250.10	-on 172.31.250.10	172.31.250.10 Communication Port: 4673
Firewall port 1234 opened between host and management server on the 192.168.250.x subnet	start -on 192.168.250.10:1234	-on 172.31.250.10:1234	172.31.250.10:1234 Communication Port: 1234

Table 122 Troubleshooting Firewalls (continued)

Configur- ation	Manual Start Parameters for CIM Extension	If Mentioned in cxws.host.parameters	Step 1 Discovery and RMI Registry Port
With 3 firewall ports opened on different ports respectively 1234, 5678, 9012.	start -on 10.250.250.10:1234 -on 172.31.250.10: 5678 -on 192.168.250.10: 9012	-on 10.250.250.10:1234 -on 172.31.250.10:5678 -on 192.168.250.10: 9012	10.250.250.1 0:1234 OR 172.31.250.1 0:5678 OR 192.168.250. 10:9012 Communication Port: 1234, 5678, 9012
With firewall port 4673 opened between host and management server. NAT environment where 10.250.250.10 subnet is translated to 172.16.10.10 when it reaches other side of the firewall	start		172.16.10.10 Communication Port: 17001

Table 122 Troubleshooting Firewalls (continued)

Configur- ation	Manual Start Parameters for CIM Extension	If Mentioned in cxws.host.parameters	Step 1 Discovery and RMI Registry Port
With firewall port 1234 opened between a host and management server. NAT environment where 10.250.250.10 subnet is translated to 172.16.10.10 when it reaches other side of the firewall	start -port 1234	-port 1234	172.16.10.10 Communication Port: 17001
With 3 firewall ports opened on different ports respectively 1234, 5678, 9012. NAT environment where all 3 NICs are translated to different 172.16.x.x subnets	start -on 10.250.250.10:1234 -on 172.31.250.10:5678 -on 192.168.250.10:9012	-on 10.250.250.10:1234 -on 172.31.250.10:5678 -on 192.168.250.10:9012	172.16.10.10: 1234 OR 172.16.20.20: 5678 OR 172.16.30.30: 9012 Communication Port: 1234, 5678, 9012
False DNS or IP is slow to resolve		jboss.properties, stop and restart service cimom.Dcxws.agency.firstwait=200000 cimom.Dcxws.agency.timeout=200000	Any IP that is reachable Communication Port: 4673

Table 122 Troubleshooting Firewalls (continued)

Configur- ation	Manual Start Parameters for CIM Extension	If Mentioned in cxws.host.parameters	Step 1 Discovery and RMI Registry Port
No DNS, never resolve		jboss.properties, stop and restart service cimom.Dcxws.agency.firstwait =200000 cimom.Dcxws.agency.timeout= 200000	Any IP that is reachable Communication Port: 4673
No firewall. Don't want to use root credentials. Want to discover with a non-existent user.	start -credentials abcuser:passwd	-credentials abcuser:passwd	Specify abcuser and password in the discovery list. Communication Port: 4673
With 3 firewall ports opened on different ports respectively 1234, 5678, 9012. Don't want to use root credentials. Want to discover with a non-existent user.	start -on10.250.250.10:1234 -on172.31.250.10:5678 -on192.168.250.10:9012 -credentials abcuser:passwd	-on10.250.250.10:1234 -on172.31.250.10:5678 -on192.168.250.10:9012 -credentials abcuser:passwd	10.250.250.1 0:1234 OR 172.31.250.1 0:5678 OR 192.168.250. 10:9012. Then, specify abcuser and passwd in the discovery list. Communication Port: 1234, 5678, 9012

Volume Names from Ambiguous Automounts Are Not Displayed

Volume names from ambiguous automounts on Solaris hosts are not displayed on the Storage Volumes page and in Capacity Manager. Some Solaris hosts have autofs and NFS mounted through an automounter. The management server cannot display volume names from ambiguous

automounts because it cannot determine if the comma separate strings that are part of the mounted volume name are host names or part of the name of a remote volume.

The following example is a comma separate string that is part of a mounted volume name. The management server cannot tell whether `test` and `three` are host names or part of the name of a remote volume. As a result, the management server does not display the volume name.

```
VolumeName = two:/ntlocal2,two:/comma,test,three,one:/ntlocal
```

Installing the Software Security Certificate

To stop receiving a Security Alert message each time you use the HTTPS logon, install the software security certificate, as described in the following steps.

IMPORTANT: Enter the DNS name of the computer in the URL instead of localhost. If you use `https://localhost` to access the management server, you are shown a “Hostname Mismatch” error.

Installing the Certificate by Using Microsoft Explorer 6.0

1. Access the management server by typing the following:
`https://machinename`
where `machinename` is the name of the management server.
2. When the security alert message appears, click **OK**.
If you do not want the Web browser to warn you about a secure connection at any Web site, select the **In the future, do not show this warning** option.
3. When you are told there is a problem with the site's security certificate, click the **View Certificate** button.
4. When you are shown the certificate information, click the **Install Certificate** button at the bottom of the screen.
5. When you are shown the Certificate Import Wizard, click **Next** to continue the installation process.
6. Select one of the following:
 - **Automatically select the certificate store based on the type of certificate** - This option places the certificate automatically in the appropriate location.
 - **Place all certificates in the following store** - This option lets you pick the store where the certificate will be stored.
7. Click **Finish**.
8. When you are asked if you want to install the certificate, click **Yes**.
You are shown the following message when the certificate is installed.

Installing the Certificate by Using Netscape Navigator 7

1. Access the management server by typing the following:
`https://machinename`
where `machinename` is the name of the management server.
2. When the security alert message appears, click the **Always** button.
3. When you are told you are requesting an encrypted page, click **OK**.
4. Click the **Always** button when you are asked if you want to accept the certificate.
5. When asked if you wanted to trust the signed applet, click the **Always** button.

Changing the Security Certificate to Match the Name of the Server

If your users are shown a Security Alert window with the following message, you might want to modify the security certificate so users feel more comfortable with installing the certificate:

"The name of the security certificate is invalid or does not match the name of the site."

You can change the security certificate so that users receive the following message instead:

"The security certificate has a valid name matching the name of the page you are trying to view."

When you change the certificate, you must use the `generateAppiqKeystore.bat` program to delete the original certificate. Then, use the `generateAppiqKeystore.bat` program to create a new certificate and to copy the new certificate to the management server.

To change the certificate:

1. Go to the `[Install_Dir]\Tools` directory, where `[Install_Dir]` is the directory into which you installed the management server.
2. To delete the original certificate, enter the following at the command prompt:
`C:\[Install_Dir]\Tools> generateAppiqKeystore.bat del`
The original certificate is deleted.
3. To create a new certificate containing the DNS name of the management server, enter the following at the command prompt:
`C:\[Install_Dir]\Tools> generateAppiqKeystore.bat`
4. If the program is unable to detect a DNS name, enter the following at the command prompt:
`C:\[Install_Dir]\Tools> generateAppiqKeystore.bat mycomputername`
where `mycomputername` is the DNS name of the computer
5. To copy the new certificate to the management server, enter the following at the command prompt:
`C:\[Install_Dir]\Tools> generateAppiqKeystore.bat copy`
The new certificate is copied to the correct location.

Troubleshooting Discovery and Discovery Data Collection

This section describes the following:

- ["Configuring E-mail Notification for Discovery Data Collection"](#) on page 511
- ["Increasing the Time-out Period and Number of Retries for Switches"](#) on page 512
- [""Connection to the Database Server Failed" Error"](#) on page 514
- ["DCOM Unable to Communicate with Computer"](#) on page 514
- ["Duplicate Listings for Brocade Switches in Same Fabric"](#) on page 514
- ["Element Logs Authentication Errors During Discovery"](#) on page 515
- ["EMC Device Masking Database Does Not Appear in Topology \(AIX Only\)"](#) on page 515
- ["Microsoft Exchange Drive Shown as a Local Drive"](#) on page 515
- ["Unable to Discover Microsoft Exchange Servers"](#) on page 515
- ["Nonexistent Oracle Instance Is Displayed"](#) on page 515
- ["Requirements for Discovering Oracle"](#) on page 515
- ["Unable to Find Elements on the Network"](#) on page 516
- ["Unable to See Path Information"](#) on page 516
- ["Device Locking Mechanism for Brocade Element Manager Query/Reconfiguration"](#) on page 516
- ["A Discovered Sun StorEdge A5000 JBOD Does Not Display Its WWN Properly"](#) on page 516
- ["Unable to Monitor McDATA Switches"](#) on page 516
- ["Unable to Detect a Host Bus Adapter"](#) on page 517
- ["Navigation Tab Displays Removed Drives as Disk Drives"](#) on page 517
- ["Unable to Obtain Information from a CLARiiON Storage System"](#) on page 517
- ["Discovery Fails Too Slowly for a Nonexistent IP Address"](#) on page 517
- [""CIM_ERR_FAILED" Message"](#) on page 518
- ["Communicating with HiCommand Device Manager Over SSL"](#) on page 519
- ["Unable to Discover a UNIX Host Because of DNS or Routing Issues"](#) on page 520

Configuring E-mail Notification for Discovery Data Collection

The management server lets you send status reports about Discovery Data Collection to users. The status reports that are sent to users can also be found in the `GAEDSummary.log` file in the `[Install_DIR]\logs` directory on the management server.

To configure the management server to send status reports on Discovery Data Collection to your e-mail account:

1. Enable e-mail notification for the management server. Refer to the User Guide for more information.
2. Add or edit the e-mail address for the Admin account.

The status reports for Discovery Data Collection automatically go the e-mail account provided for the Admin user. To add or edit an e-mail address for the Admin account, log in as Admin and then follow the steps in “[Modifying Your User Profile](#)” on page 80.

3. If you want additional users to receive the status reports for Discovery Data Collection, do the following:
 - a. Click **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
 - b. Click **Show Default Properties** at the bottom of the page.
 - c. Copy the `gaedemail` property. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
 - d. Return to the Advanced page (**Options > Storage Essentials > Manage Product Health**). Then, click **Advanced** in the Disk Space tree).
 - e. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
 - f. Assign the e-mail accounts you want to receive the report to the `gaedemail` property. For example, if you want `user1@appiq.com` and `user2@appiq.com` to receive these status reports, modify the `gaedemail` property in the **Custom Properties** field as follows:
`gaedemail=user1@appiq.com;user2@appiq.com`

NOTE: Make sure the hash (#) symbol is removed from the `gaedmail` property.

- g. When you are done, click **Save**.
 - h. Restart the service for the management server for your changes to take effect.

Increasing the Time-out Period and Number of Retries for Switches

If you are having difficulty obtaining information from your switches during Discovery Data Collection, you may need to increase the time-out period and the number of retries. By default, the management server gives a switches five seconds to respond to its requests for information during Discovery Data Collection. If the switch does not respond the first time, the management server tries again. The management server says it cannot contact the switch if it does not receive a response from the switch a second time.

To change the time-out period and number of retries for switches, modify the properties specified [Table 123](#) on page 513 and [Table 124](#) on page 513 as described in the following steps:

1. Access the management server.
2. Select **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
3. Click **Show Default Properties** at the bottom of the page.

4. Copy the commands specified in [Table 123](#) on page 513. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
5. Return to the Advanced page (**Options > Storage Essentials > Manage Product Health**). Then, click **Advanced** in the Disk Space tree).
6. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
7. Make sure the property is not commented out by removing the hash (#) symbol in front of the property. To modify the time-out period, set the corresponding property for your switch in the following table to the number of millisecond you want. The default is 5000 ms. For example, to change the time-out period to 30000 ms for a McDATA switch, you would set the `cimom.McData.Snmp.Timeout` property to 30000, as shown in the following example:
`cimom.McData.Snmp.Timeout=30000`

Table 123 Time-out Properties

Switch	Property
McDATA/Connectrix discovered through SNMP	<code>cimom.McData.Snmp.Timeout</code>
Cisco	<code>cimom.Cisco.Snmp.Timeout</code>
Other switches discovered through SNMP: <ul style="list-style-type: none"> • CNT • Sun StorEdge • QLogic 	<code>cimom.snmp.switch.timeout</code>

8. To modify the number of retries, repeat Steps 4 through 6 by copying and pasting the property specified in [Table 124](#) on page 513. Set the corresponding property for your switch in the following table to the number of retries you want. The default is two retries. For example, to change the number of retries to five for a McDATA switch, set the `cimom.McData.Snmp.Retries` properties as shown in the following example:
`cimom.McData.Snmp.Retries=5`

Table 124 Retry Properties

Switch	Property
McDATA/Connectrix discovered through SNMP	<code>cimom.McData.Snmp.Retries</code>
Cisco	<code>cimom.Cisco.Snmp.Retries</code>

Table 124 Retry Properties (continued)

Switch	Property
Other switches discovered through SNMP: <ul style="list-style-type: none">• CNT• Sun StorEdge• QLogic	<code>cimom.snmp.switch.retries</code>

9. When you are done, click **Save**.

10. Restart the service for the management server for your changes to take effect.

While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

“Connection to the Database Server Failed” Error

If you received an error message resembling the following after getting all element details, verify that the database instance is running:

```
The connection to the database server failed. Check that the Oracle instance 'OIQ3' on host '192.168.1.162:1521' is running correctly and has the management software for Oracle installed correctly.
```

Assume you received the error message listed above. You would want to verify the following:

- Oracle instance OIQ3 on host 192.168.1.162 port 1521 is running.
- The management software for Oracle is installed on the server running the Oracle instance. One of the installation's tasks is to create an APPIQ_USER user account with enough privileges for the software to view statistics from the database.

Once you have verified the items listed above, run “Discovery Data Collection” again. If you continue to see the error message, contact customer support.

DCOM Unable to Communicate with Computer

Sometimes the following error message appears in the event log of the management server when the software is monitoring a Brocade switch:

```
DCOM was unable to communicate with the computer 192.168.10.21 using any of the configured protocols
```

where 192.168.10.21 is the IP address of the Brocade switch.

Ignore this error message.

Duplicate Listings for Brocade Switches in Same Fabric

If you discover more than one Brocade switch in the same fabric, the **Targets** tab displays duplicate listings for the Brocade switches. Each Brocade switch is listed multiple times with the IP address of the other switches and its own.

For example, assume you discovered Brocade switches QBrocade2 and QBrocade5 in the same fabric, the switches are listed twice on the **Targets** tab. QBrocade2 is listed twice, once with its own IP address, the other time with the IP address of QBrocade5, as shown below:

```
192.168.10.22 Switch QBrocade2, QBrocade5 admin
192.168.10.25 Switch QBrocade2, QBrocade5 admin
```

Element Logs Authentication Errors During Discovery

During discovery, you may see SNMP authentication errors on the element you are trying to discover. The management server is probing the element with an SNMP request. If the element does not know the management server, it logs authentication errors.

EMC Device Masking Database Does Not Appear in Topology (AIX Only)

An EMC device masking database attached to an AIX host does not appear in the Topology tree under the **Application Path - Unmounted** node on the **Topology** tab in System Manager.

If the EMC device masking database is attached to a host running Microsoft Windows or Sun Solaris, the masking database appears under the **Application Path - Unmounted** node.

Microsoft Exchange Drive Shown as a Local Drive

Microsoft Exchange Servers have a drive M. The software displays this drive as a local fixed disk, instead of a Microsoft Exchange Server special drive.

Unable to Discover Microsoft Exchange Servers

If DNS records for your Microsoft Exchange servers are outdated or missing, the discovery of Microsoft Exchange may fail because Microsoft Exchange is dependant on Active Directory, which is dependant on DNS. Since Active Directory is dependant on DNS, Active Directory replication and Active Directory lookups may fail or contain errors if DNS records are not accurate.

Nonexistent Oracle Instance Is Displayed

The software uses the Oracle Transparent Name Substrate (TNS) listener port to detect Oracle instances on a server. Sometimes an Oracle instance is removed from the server, but not from the TNS listener port. This results in the software detecting the nonexistent Oracle instance and displaying it in the topology. Refer to Oracle documentation for information on how to remove the deleted Oracle instance from the TNS listener port.

Requirements for Discovering Oracle

To discover Oracle:

- The management software for Oracle must be installed. For information about installing the management software for Oracle, refer to the *Installation Guide*.
- By default, the software sets the TNS Listener Port to 1521. If you use another port, you can change the port number on the Discovery Targets tab.
- Oracle discovery relies on the TNS networking substrate on which Oracle is built (TNS is Oracle's proprietary protocol). The software does not use TNS listener password. If you have

set a TNS Listener password, the software is not able to discover the Oracle instances serviced by the listener.

Unable to Find Elements on the Network

The management server uses ping to find the devices on the network enabled for IP. Ping is a program that lets you verify that a particular IP address exists. Ping is not guaranteed to return a response from all devices. If Discovery is not able to find a device automatically, enter the IP address for the device on the Discovery Targets tab, which can be accessed by clicking the Discovery button at the top of the screen in the management server. Sometimes ping cannot find the device if one of the following conditions occur:

- Network configuration does not support ping, including data center security (firewalls).
- Device has the ping responder turned off.
- Device does not support ping.

Unable to See Path Information

You will not be able to see path information if LUN masking information is missing. To view LUN masking information, follow the steps described in the section, "[Accessing Information About Host Security Groups](#)" on page 275.

Device Locking Mechanism for Brocade Element Manager Query/Reconfiguration

Please keep in mind that the configuration for Brocade switches is locked while getting all details for elements in a zones. The software ensures that each CIM query locks out any reconfiguration. For example, if you are getting details for elements in all zones, you cannot add a new Brocade switch while your doing it (the discovery or configuration process waits until the collection of details is finished before proceeding). However, simultaneous CIM queries do not lock each other out.

A Discovered Sun StorEdge A5000 JBOD Does Not Display Its WWN Properly

Although full monitoring and management support is available only to those devices for which there is a provider, the software's topology displays other devices found on your storage area network (SAN) to give you a more complete view. However, because these devices do not have a provider, only basic information is returned. In some cases, as with the Sun StorEdge A5000 JBOD (Just a Bunch of Disks), the Worldwide Name (WWN) presented and reported to the management server may be different from the official WWN of the device, as the management server reports the WWN of the port connected to the fabric.

Unable to Monitor McDATA Switches

McDATA switches use the Fibre Channel Switch Application Programming Interface (SWAPI) to communicate with devices on the network. The McDATA switches allow only one SWAPI connection at a time. For example, if the management server discovers the IP address of the McDATA switch, other management servers and third-party software are not able to communicate with the switch using SWAPI.

Use Enterprise Fabric Connectivity (EFC) Manager to communicate with the McDATA switch. EFC Manager versions 7.0 and later can communicate with the management server and the switch. This configuration lets multiple instances of the management server or other clients contact EFC Manager, which in turn provides information about the switch. To communicate with the EFC Manager, first install the Bridge Agent. Then, enter the IP address of the server running EFC Manager in the Discovery pane. The user name and password must be for EFC Manager, and the user name and password are case sensitive. Refer to your McDATA representative for more information about the Bridge Agent and EFC Manager.

IMPORTANT: EFC Manager uses the SWAPI connection, preventing other third-party software from contacting the switch.

Unable to Detect a Host Bus Adapter

The software is unable to detect a host bus adapter if you install its driver before you have completed installing the Solaris operating system for the first time. For example, you installed the HBA drives too early when you used JumpStart to install Solaris. The best way to install the HBA driver is to install it after Solaris has been installed and is running.

Navigation Tab Displays Removed Drives as Disk Drives

If you remove an internal disk from a Solaris host and do not enter the `cfgadm` command, the Navigation tab displays the empty slot as `DiskDrives_XXXXX` after getting element details. The `cfgadmn` command makes the software realize the drive has been removed. Refer to the documentation that shipped with the Solaris operating system for more information about the `cfgadm` command.

Unable to Obtain Information from a CLARiiON Storage System

If you are having difficulty obtaining topology information or element details from a CLARiiON storage system, the NaviCLI might have timed out as a result of the service processor being under a heavy load. The management server uses the NaviCLI to communicate with the CLARiiON storage system. This situation has been seen in the field when the service processor is running more than 35,000 IOPS (IOs/Sec).

Try obtaining Discovery Data Collection from a CLARiiON storage system when the service processor is not under such a heavy load.

Discovery Fails Too Slowly for a Nonexistent IP Address

If you enter a nonexistent IP address, the management server times out by default after 20 seconds on Windows and after three minutes on Linux. If you want to shorten the time-out period, modify the `cimom.CimXmlClientHttpConnectTimeout` property as described in this section.

NOTE: The management server does not accept a period longer than its default setting. If you set `cimom.CimXmlClientHttpConnectTimeout` property to more than 20 seconds on Windows, the management server ignores the values of this property and reverts back to the default settings.

To modify the default time-out:

1. Access the management server.
2. Click **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
3. Click **Show Default Properties** at the bottom of the page.
4. Copy the `cimom.CimXmlClientHttpConnectTimeout` property you want to modify. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
5. Return to the Advanced page (**Options > Storage Essentials > Manage Product Health**). Then, click **Advanced** in the Disk Space tree).
6. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
7. Make your changes in the **Custom Properties** field. Make sure the property is not commented out by removing the hash (#) symbol in front of the property.
8. To modify the time-out period, set the `cimom.CimXmlClientHttpConnectTimeout` property to the number of millisecond you want. For example, to change the time-out period to 200 ms, set the `cimom.CimXmlClientHttpConnectTimeout` property, as shown in the following example:

```
cimom.CimXmlClientHttpConnectTimeout=200
```
9. When you are done, click **Save**.
10. Restart the service for the management server for your changes to take effect.

While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

"CIM_ERR_FAILED" Message

If you are in a McDATA environment where the EFC Manager Service Processor is managing multiple switches, it is possible that the management server will send SWAPI requests faster than the EFC Manager Service Processor can handle them. The management server may detect this as a failed connection and take corrective action. When this happens, you are shown a "CIM_ERR_FAILED" message whenever the management server tried to access the McDATA switches and directors.

The management server then attempts to reconnect to the EFCM by creating a new SWAPI connection. EFCM versions 8.x and later have five SWAPI connections available. EFCM versions 7.1.3 and later but before version 8.x have three SWAPI connections available. If the management server reconnects successfully, a reconnect event is generated and no further action is necessary.

If the management server cannot reconnect to the EFCM, another event is generated with a severity of major. If this happens, any operation the management server is performing (Discovery Data Collection) involving switches on that EFCM fails.

To prevent the “CIM_ERR_FAILED” messages, increase the delay between the management server’s SWAPI calls to EFCM, as described in the following steps:

1. Click **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy `cimom.mcData.swapiThrottle=200`. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
4. Return to the Advanced page (**Options > Storage Essentials > Manage Product Health**). Then, click **Advanced** in the Disk Space tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make your changes in the **Custom Properties** field by changing the value of `cimom.mcData.swapiThrottle`. For example, the default is 200 milliseconds. To change the value to 800 milliseconds, change the xxx value to 800, as shown in the following example:

```
cimom.mcData.swapiThrottle=800
```

NOTE: If you want no delay, change the value to 0 for 0 milliseconds. The maximum delay you can have is 1,000 milliseconds (`cimom.mcData.swapiThrottle=1000`),

7. When you are done, click **Save**.
8. Restart the service for the management server for your changes to take effect.
While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Communicating with HiCommand Device Manager Over SSL

By default the management server communicates with HiCommand Device Manager through a nonsecure connection. You can configure the management server so that it communicates with HiCommand Device Manager over a secure socket layer (SSL) connection by doing one of the following:

- **Use HTTPS in the discovery address** - Prepend `https://` to the discovery address to force the connection to HTTPS mode, for example, `https://192.168.1.1`, where 192.168.1.1 is the IP address of the host running HiCommand Device Manager. Use this option if you have one HiCommand Device Manager you want to communicate through a secure connection (SSL) and another you want to communicate through a nonsecure connection.
- **Modify an internal property** - Change the value of the `cimom.provider.hds.useSecureConnection` to true, as described in the steps in this

section. Use this option if you want all connections to HiCommand Device Manager to be secure (SSL).

To set all connections with HiCommand Device Manager to SSL:

1. Click **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the `cimom.provider.hds.useSecureConnection` property. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text and then right-click the selected text. Then, select **Copy**.
4. Return to the Advanced page (**Options > Storage Essentials > Manage Product Health**). Then, click **Advanced** in the Disk Space tree).
5. Paste the copied text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.
6. Make your changes in the **Custom Properties** field. Make sure the property is not commented out by removing the hash (#) symbol in front of the property.
7. Change the value assigned to the `cimom.provider.hds.useSecureConnection` property to true, as shown in the following example:
`cimom.provider.hds.useSecureConnection=true`
8. When you are done, click **Save**.

If you want to connect to another instance of HiCommand Device Manager by using a nonsecure connection, prepend `http://` to the discovery address to force the connection to nonsecure mode, for example, `http://192.168.1.1`, where 192.168.1.1 is the IP address of the host running HiCommand Device Manager.

9. Restart the service for the management server for your changes to take effect.

While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Important: While the AppStorManager service is stopped, the following occurs:

- Users are not be able to access the management server.
- The management server is unable to monitor elements at this time.

Unable to Discover a UNIX Host Because of DNS or Routing Issues

If the management server is unable to discover a UNIX host because of a DNS or routing issues, you will need increase the amount of time that passes before the management server times out for that CIM Extension. By default, the management server waits 1,000 ms before it times out. It is recommended you increasing the time before the management server times out to 200000 ms (3.33 minutes), as described in the following steps. If you continue to see time out issues, you can still increase the time before the management server times out, but keep in mind that it will lengthen discovery.

To increase the time out period:

1. Select **Options > Storage Essentials > Manage Product Health**. Then, click **Advanced** in the Disk Space tree in the management server.
2. Paste the following text into the **Custom Properties** field. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the field and select **Paste**.

```
cimom.cxws.agency.firstwait=200000  
cimom.cxws.agency.timeout=200000
```

where

- `cimom.cxws.agency.firstwait` - The `firstwait` property controls the amount of time required for the management server to wait after it first contacts the CIM Extension on the host before the management server attempts to proceed with a username and password. The default value is 1,000 ms. You are modifying it to wait 20,000 ms or 3.33 minutes.
 - `cimom.cxws.agency.timeout` - The `timeout` property controls the allowable interval of silence before either the CIM Extension or the management server start to question whether its partner is still alive. If an entity (management server or extension) has not received a message from the other during the interval set by the `timeout` property, it will send an "are you there" message. If that message is not acknowledged during the interval set by the `timeout` property, the entity will conclude that the connection is no longer functioning. The CIM Extension will stop attempting to make a connection. When this occurs on the side of the management server, the management server will attempt to re-connect (and it will keep re-attempting until the host becomes available). The default value is 1,000 ms. You are modifying it to wait 20,000 ms or 3.33 minutes.
3. Click **Save**.
 4. Restart the service for the management server for your changes to take effect.
While AppStorManager is restarting, users are not able to access the management server. The AppStorManager service must be running for the management server to monitor elements.

Troubleshooting the Java Plug-in

This section describes the following:

- [s"Java Applet Has Data from a Different Version of Management Server Software"](#) on page 521
- ["OutOfMemoryException Messages"](#) on page 521
- ["Improving Reload Performance in System Manager"](#) on page 522

Java Applet Has Data from a Different Version of Management Server Software

If you attempt to monitor a host with old JAR (Java Archive) files, you might be unable to monitor the host and you might see the following error message:

```
The Java applet has data from a different version of the management server.  
Please close and re-start your browser.
```

The reason for this error message is that the client still has JARs from the previous build in its Java Plug-in cache. To remove the old JARs, clear the cache for the Java plug-in.

OutOfMemoryException Messages

In some rare cases it may be necessary to increase the amount of memory for the Java plug-in on the client computer. This should only be done if you are seeing `OutOfMemoryException` messages in the Java console on the client side.

Improving Reload Performance in System Manager

If your Java plug-in control panel cache setting is set at 50 MB, it is recommended you increase this setting to 150 MB or more. Increasing this setting improves the reloading performance of System Manager.

Troubleshooting Provisioning

This section describes the following:

- ["Cannot Access a Resource Owned by Another Controller"](#) on page 522
- ["Error -56"](#) on page 522
- [""Can't delete this zone" Message"](#) on page 522
- ["Changes in EFC Manager Requiring "Discovery Data Collection""](#) on page 523

Cannot Access a Resource Owned by Another Controller

You received the message about not being able to access a resource owned by another controller because you tried to access a controller that has not been discovered. You should discover all controllers on the Engenio storage system.

For example, assume you discovered only one of the controllers on an Engenio storage system with two controllers. If you want to change a volume, such as add or delete a LUN, you will not be able to make the change to the volume associated with the controller that has not been discovered.

To discover the controller, do the following:

1. Click the **Discovery** menu.
2. Enter the IP address for the controller on the **Discovery Targets** tab.
3. Click the **Discover Selected** button.
4. Click the **Discovered Elements** tab.
5. Click the **Get Details** button.
6. See ["Discovering Filers, Tape Libraries, Switches, and Storage Systems"](#) on page 13 for more information.

Error -56

If you see `error -56`, the switch has network connection failures/problems. To solve the problem, make sure the switch is physically connected to the network. Then, redo the task you were originally trying to complete.

If you now see `-21 (OBJECT_NOT_FOUND)` errors, the switch needs to be rediscovered.

“Can't delete this zone” Message

If you are shown the following message when you try to delete a zone, move the zone to an inactive zone set. Then, delete the zone.

Can't delete this zone, it is member of an Active Zoneset

Changes in EFC Manager Requiring “Discovery Data Collection”

If you use EFC Manager to delete zones or zone sets, perform “Discovery Data Collection” (**Discovery > Details**) on the management server afterwards. The changes are not reflected by the management server until “Discovery Data Collection” is done.

Troubleshooting Hardware

This section describes the following:

- “About Swapping Host Bus Adapters” on page 523
- ““Fork Function Failed” Message on AIX Hosts” on page 523
- “Known Driver Issues” on page 523
- “Known Device Issues” on page 524
- ““mailbox command 17 failure status FFF7” Message” on page 527
- ““Process Has an Exclusive Lock” Message” on page 527

About Swapping Host Bus Adapters

Swapping brands of host bus adapters (HBA) on a Microsoft Windows 2000 host may have undesirable side effects. For example, after swapping out one brand of an HBA for another (including driver installation), `WinMgmt.exe` might crash repeatedly and appear to be associated with an error in the Windows Event Log about being unable to retrieve data from the `PerfLib` subkey in the Registry. To solve this problem, reinstall the operating system.

“Fork Function Failed” Message on AIX Hosts

If a CIM Extension running on AIX detects low physical or virtual memory while starting, a “Fork Function Failed” message appears. A CIM Extension on AIX uses additional memory and CPU resources at start time. If the resources on the AIX machine is already low, you may see the “Fork Function Failed” message. Depending on the AIX operating system or hardware, the host may crash after you see this message.

Known Driver Issues

If you are having problems with a driver, keep in mind the following:

- The software requires the driver to have a compliant SNIA HBA API. Emulex driver version 4.21e does not support the SNIA HBA API.
- If the driver has a compliant SNIA HBA API, make sure the driver is installed correctly.

Known Device Issues

The following table provides a description of the known device issues. You can find the latest information about device issues in the release notes.

Table 125 Known Device Issues

Device	Software	Description
AIX host	NA	<p>If you are receiving replication errors for an AIX host, the provider may be trying to connect to the host using the 0.0.0.0 IP address instead of the real host IP address. If this situation is occurring, you would see a message containing the following when you start the CIM Extension:</p> <pre>CXWS 3.1.0.144 on 0.0.0.0/0.0.0.0 now accepting connections</pre> <p>To fix this situation, add the following line to the /opt/APPQcime/tools/start file on the AIX host:</p> <pre>export NSORDER=local,bind</pre>
AIX host using an IBM Storage System	NA	<p>If you have an AIX host using an IBM storage system, not all bindings may be displayed on the bindings page on the Navigation tab. For example, assume diskA on host123 has six paths. All six bindings may not be displayed.</p>
Hosts running SGI IRIX version 6.5.22 or 6.5.24	NA	<p>If a host running SGI IRIX version 6.5.22 or 6.5.24, the HBA port page on the Navigation tab in System Manager displays 0 GB/s for HBA ports.</p>

Table 125 Known Device Issues (continued)

Device	Software	Description
SGI IRIX host	CXFS file systems	<p>The management server can only monitor CXFS file systems from the host generating the input/output. For example, assume the elements are part of a CXFS file system. When you generate input/out into the metadata server into /folder, only the metadata server is able to monitor the file system. For example, assume the metadata server generates 100 KB write, the management server displays 0 KB write for /folder on the metadata client.</p> <p>See "About the Data from CXFS File Systems" on page 238 for more information.</p>
Solaris host	Sun SAN Foundation Suite driver (Leadville driver)	The bindings page reports a SCSI number that comes from the HBAAPI. This number cannot be seen by the user. For example SCSI target 267008 does not correlate to anything.
Solaris host	HDLM	<p>If you sync the Solaris host by itself without the switches and storage, the storage volume page reports all drive types as local.</p> <p>Once you discover the host with the switches and storage, it reports its drives as being external. It was the same result with Active-Active and Active-Standby.</p>
Solaris host	HDLM	Solaris HDLM disks cannot be monitored. If you try monitoring them, the management server displays a message saying "data is late or an error occurred."
Solaris host	HDLM	<p>Do a Discovery Data Collection for the host by itself. In the bindings page, the controller number are displayed as c-1. For example c-1t0d58.</p> <p>Perform Discovery Data Collection on the host with storage and switches. The controller numbers are displayed correctly.</p>

Table 125 Known Device Issues (continued)

Device	Software	Description
Solaris host	VxVM	<p>If you discover a host with any typical SAN disk groups off line, the storage volume page shows SAN mount points as local instead of external. These disks, however, are not accessible.</p> <p>When you perform Discovery Data Collection with all disk groups online, disks on the SAN are shown as external. Hosts connected directly to a storage system are shown as local, except for hosts connected by fibre. Hosts connected directly to a storage system through fibre are shown as external.</p>
Windows host	VxVM	<p>When a Windows host with VxVM is used, the SCSI bus number is always reported to be one in the SCSI bus column of the Disk Drives page.</p>
Any host	NA	<p>The Unmounted Volume field under Capacity Summary automatically displays 0 MB if you discovered the host but not the storage system connected to it. This may occur if you did not enter the IP address of the storage system when performing discovery and/or your license does not allow you to discover a particular storage system. See the Supported Elements section in the "List of Features" to determine which storage systems you can discover. The "List of Features" is accessible from the Documentation Center (Help > Documentation Center in Storage Essentials).</p>

Table 125 Known Device Issues (continued)

Device	Software	Description
IBM Storage Systems	Subsystem Device Driver (SDD) or MPIO (Multi Pathing IO)	If you discovery an IBM storage system without SDD, incorrect stitching is displayed in System Manager for the storage system. You are shown only one path if the storage system is using MPIO instead of SDD.

“mailbox command 17 failure status FFF7” Message

If one or more of your Microsoft Windows hosts are using an Emulex HBA driver, you may see the following message in Windows Event Viewer:

“mailbox command 17 failure status FFF7”

This message can be safely ignored. The HBA API is being used to access data in the FLASH memory of the adapter that does not exist and this is causing the event to be logged. This issue has been seen with version 5.2.2 of the driver.

“Process Has an Exclusive Lock” Message

You will receive a message resembling the one shown below if a process has already locked the EMC Symmetrix storage system and you attempt a process that requires a lock on the Symmetrix storage system. The Symmetrix storage system can become locked for many reasons. For example, the storage system becomes locked when it performs LUN mapping, LUN masking or Discovery Data Collection. The Symmetrix storage system may also remain locked after a provisioning operation has failed.

“SYMAPI routine SymDevMaskSessionStart failed with error code 188: The operation failed because another process has an exclusive lock on the local Symmetrix.”

After the management server has detected the lock on the Symmetrix storage system, it tries to access the storage system for 15 minutes and logs the errors.

If you receive the error message, determine if someone is performing an operation that requires a lock, such as LUN mapping, LUN masking or Discovery Data Collection. This also applies even if one of the processes is being used by a third-party product, such as for LUN masking. If so, wait until the process is complete. Only manually remove the lock if you are certain that no other processes are occurring on the storage system. To learn how to remove the lock, refer to the documentation for the Symmetrix storage system.

If a provisioning failure has caused the Symmetrix storage system to remain locked, you are alerted to this situation in Event Monitoring for Storage Essentials and on the Properties tab. You may receive a message resembling the following:

```
Unable to end device masking session. Symmetrix '000001835005700' may be
locked.
```

Glossary

A

- access point** It is the intersection of the IP address and the provider that discovered the IP address. It is displayed on the screens for discovery. A provider is software that is used to gather information about an element.
- active zone set** An active zone set is the zone set in use. You can have only one zone set active at a time; however, you can have a zone in more than one zone set. Zones sets are usually created for a particular task.
- Zones work by exclusion. Members of a zone can communicate with other members in the zone. If two devices are not within the same zone, they cannot communicate. Only elements in active zones can communicate with each other. When a zone is not active, it does not have any effect.
-

C

Common Information Model (CIM)

The Common Information Model is a common data model of an implementation-neutral schema for describing overall management information in a network/enterprise environment. CIM is comprised of a specification and a schema. The specification defines the details for integration with other management models (i.e. SNMP's MIBs or the DMTF's MIFs) while the schema provides the actual model descriptions.

Common Information Model Object Manager (CIM Object Manager)

A component in the CIM management infrastructure that handles the interaction between management applications and providers. The CIM Object Manager supports services such as event notification, remote access, and query processing. The CIM Object Manager also grants access to the CIM Object Manager repository.

F

File Server Storage Resource Management (SRM)

File Server Storage Resource Management (SRM) does a recursive lookup on the file system and stores the information in an embedded database. File System Viewer can

scan files very quickly because of its structure in the database and because it uses a multi-threaded process. More than one process can be used at a time to scan the files.

D

device This documentation set defines a device as a piece of hardware in the storage network.
EMC uses the term device to refer to a volume on one of its storage systems.

E

element An element is anything on the network that can be detected by the management server, such as hosts and switches.

element created in Chargeback Manager An element created in Chargeback Manager is a type of generic element. When you create a record for an element in Chargeback Manager, the element appears as a generic element in the topology.

G

generic element An element is considered to be generic if the management server can detect the element but it cannot obtain additional information about the element during “Get the Topology” or “Discovery Data Collection.”

global reporting view A global reporting view contains information in the database that can be used for global reports.

Global Reporter server A management server that has global reporting enabled.

H

host persistent binding A system SCSI target ID assigned permanently to an element. The host binding is implemented on the host bus adapter (HBA), resulting in the HBA being tied to a certain LUN.

I

initiator WWN The Worldwide Name (WWN) of a host bus adapter's port. The WWN differentiates the port from others.

inode file An inode file stores information about a file, excluding the file's data.

M

Managed Application

Licenses (MALs) Managed application licenses (MALs) are the number of detected instances of Microsoft Exchange, Oracle and Sybase Adaptive Server Enterprise.

managed object A hardware or software system component that is represented as an instance of the CIM class. Information about managed objects is supplied by data and event providers, as well as by the CIM Object Manager.

Managed Access Points (MAPs)

Manage access points (MAPs) are the sum of all storage access ports of all hardware elements that the management server manages.

materialized view A materialized view is a snapshot of data, from the database, created from a query.

mapped Mapped is capacity that is accessible by one or more hosts external to the array (aggregated capacity of volumes that are accessible from hosts external to the subsystem).

meta device This term is used by EMC. A meta device is a device that is a concatenation of several devices.

metavolume Metavolumes are created from a disk, slice, stripe, or other metavolumes. Metavolumes are extremely useful because they can expand their storage capacity, such as to mainframe volume sizes. Also referred to "LDEVs" for HDS storage systems.

missing element The management server was able to discover the element, but it lost contact with the element before more information could be gathered during "Get the Topology" or "Discovery Data Collection." A missing element can be managed if the management server lost contact with the element after Discovery Data Collection was performed.

multipathing The process of providing a server more than one path to a storage system. So that in the case of an emergency, the server will have continuous access to the storage system. Multipathing can be done many ways. For example, you can provide redundant switches for a server to access a storage system. Another example of multipathing is providing redundant paths from the server to the switch.

P

provider A provider is software that is used to gather information about an element.

proxy server A device, such as a host, connected to a storage system. It is sometimes referred to as a storage system proxy or an API proxy. An example of a proxy server is the EMC

Solutions Enabler or Hitachi HiCommand Device Manager. Engenio storage systems do not require a proxy, as they can be accessed directly.

S

SAN

A Storage Area Network (SAN) is a high speed network configuration that is dedicated to transporting storage data among network devices, such as storage systems, hosts (servers), switches, and tape libraries to end users. In addition to connecting local elements to storage arrays, it may also be extended to off site or remote locations for the purposes of backup, archival or acting as a hot site in the event of a disaster.

A SAN can communicate via current technologies such as ESCON (mainframe), fibre channel, or newer technology such as iSCSI. SAN's can support several configurations such as disk mirroring, RAID 5, backup/restore, and data migration, as well as being able to incorporate Network Attached Storage (NAS).

SMTP

Simple Mail Transfer Protocol (SMTP) is a TCP/IP protocol used in sending and receiving e-mail.

Storage Management Infrastructure Specification

A Storage Networking Industry Association (SNIA) standard for implementing data storage management using the Common Information Model (CIM).

storage pool

It is a group of volumes. Also known as volume group.

switch port zoning

A type of zoning in which the port of the switch is physically in the zone. Any device attached to the port is automatically in the zone.

V

virtual application

A placeholder you create for an unsupported application. For example, assume your company has created an internal application, and you want to be able to use the software to keep track of that application. You can create a virtual application for that product.

W

Web-Based Enterprise Management (WBEM)

Web-Based Enterprise Management (WBEM) is an initiative based on a set of management and Internet standard technologies developed to unify the management of enterprise computing environments. WBEM provides the ability for the industry to deliver a well-integrated set of standard-based management tools leveraging the emerging technologies such as CIM and XML.

Windows Management

Instrumentation

(WMI)

Microsoft created WMI as its implementation of Web-based Enterprise Management (WBEM). For more information about WMI, refer to the Microsoft Web site at <http://www.microsoft.com>.

The Windows CIM Extension enhances Windows Management Instrumentation (WMI) so that it can gather information from host bus adapters and make the information available to the management server.

Worldwide Name

(WWN) zoning

A type of zoning in which the port is assigned to a Worldwide Name of a host or a storage system. It is not dependent on the switch.

WWN (Worldwide

Name)

Fibre-channel Worldwide Name. Usually formatted as 16-hexadecimal digits. This name is globally unique, and it identifies the connection or set of connections to the network.

Z

zone

A collection of zone aliases and ports.

zone alias

To avoid remembering a port's Worldwide Name (WWN), assign the port to a zone alias.

zone member

A zone member is a port attached to a switch, a Worldwide Name (WWN) or a zone alias.

As a best practice, a zone should contain either zone aliases or ports, but not both.

zone set

A zone set is a collection of zones. You can have only one zone set active in a fabric; however, you can have a zone in more than one zone set. Zones sets are usually created for a particular task.

Index

A

about

- Access tab 169
- asset attributes 234
- Business Tools 495
- buttons in System Manager 164
- Capacity Manager 423
- Chargeback Manager reports 484
- custom commands 204
- custom reports 350
- CXFS 238
- Event Manager 323
- Events tab 233
- filtering events 333
- groups 197
- Home page 6
- List tab 167
- monitoring options 408
- Navigation screen 215
- Path tab 175
- Performance Manager 400
- Policies tab 237
- Policy Manager 437
- provisioning 239
- Provisioning tab 233
- remote console 211
- Reporter 337
- Reports tab 236
- security 71
- storage tiers 458
- System Manager 161
- System Manager user interface 164
- Topology Screen 223
- virtual applications 532
- zone aliases 533
- zone sets 533
- access point 529
- Access tab 164, 169
- accessing
 - asset attributes 461
 - Database Admin Utility 150
 - domain controller 51
 - e-mail schedules 349
 - events 335
 - host security groups 275
 - logs 106
 - Navigation screen 219
 - Policy Manager 438
 - properties 219
 - storage pools 262
 - System Manager 163
 - trending information 435
 - volume information 263
 - zone sets 249
 - zoning information 247
- accessing information
 - zone aliases 244
- account
 - password 79
- Accounters 495
- accounts
 - users 78
- activating
 - a zone set 253
 - zone sets 249
- active
 - zone sets 249
- Active Client Logons 408
- Active Directory 515
- active zone set 529
- adding
 - asset record 459
- adding
 - asset information 463
 - custom command 204
 - custom information 466
 - departments 467
 - discovery policies 441
 - discovery schedule 100
 - domain controller 51, 66
 - elements 86, 88
 - e-mail schedule 344
 - event policies 444

- general information 464
- geographic information 466
- host security groups 276
- hosts 302, 311
- IP range 15
- license 96
- licensing information 466
- organizations 86
- provisioning policies 442
- roles 83
- staff information 465
- storage pool 261
- storage volume 265
- switches 44
- tape drive 222
- TNS Listener Port 66
- user accounts 78
- utilization policies 438
- virtual application 176
- virtual applications 232
- warranty information 466
- zone alias members 245
- zone members 248
- zone sets 249
- zones 247
- Advisors 495
- aging statistics 126
- AIX 515
- API proxy 531
- appiq.log 106, 108
- APPIQ_OWNER account 51
- APPIQ_USER 67
- application
 - virtual 176
- Application Administrator role 71
- applications
 - business costs 192
 - discovering 51
 - virtual 232
- archive destination 157
- archive mode
 - setting 155
- arranging
 - elements 185
- asset attributes
 - elements 234
 - viewing 461

- asset based
 - reports 484
- asset information
 - adding 463
- Asset Management 495
- asset record
 - creating 459
- asset-based chargeback
 - setting up 468
- assigning
 - business costs 192
 - costs 458
 - Custom Name 221
 - e-mail schedules 344
 - MIME types 44, 69
- audience xix
- authentication errors
 - SNMP 515
- authorized reseller, HP xxi
- automatically
 - clearing events 111
 - deleting events 111
- Automators 495
- available pools 262
- Average Blocks Requested 408
- Average Delivery Time 408

B

- backing up
 - database 148, 150
- bar chart 401
- binding
 - persistent 530
- block size
 - cache 287
- Bridge Agent 18
- Brocade
 - switches 331
 - tracing 108
- Brocade Rapid program 37
- Brocade switches 37
 - discovering 15
 - events 332
- browsing
 - hosts 214
 - switches 214
- Buffer Hit Ratio 408

- business cost
 - setting 176
- business costs
 - assigning 192
- Business Tools 495
 - risk Analysis 499
- buttons
 - remote console 211, 212
 - right 428
- Bytes Received 408
- Bytes Transmitted 408
- C
- cache
 - refreshing reports 128
- cache block size
 - changing 287
- cache read-ahead multiplier 260
- cache settings
 - modifying 260, 287
- calculating
 - fixed declining balance 478
 - straight line depreciation 477
- calculations for
 - double declining balance 479
- capacity charts 433
- Capacity Manager
 - about 423
 - capacity charts 433
 - Capacity tab 423
 - finding capacity 428
 - global view 425
 - HDS 428
 - legend 423
 - magnifying 425
 - Path tab 423
 - printing 425, 431
 - toolbar 425
 - trending 435
 - utilization reports 431
 - Utilization tab 423
 - utilizing reports 431
- Capacity tab 423
- categorizations
 - storage 458
- certificate
 - installing 9

- changing
 - archive destination 157
 - cache block size 287
 - cache settings 260
 - custom commands 207
 - database 45
 - database passwords 151
 - date 123, 139
 - discovery policies 448
 - domain controller 51, 66
 - e-mail address 80
 - e-mail schedules 125
 - event policies 451
 - external tools 214
 - fabric name 176, 200
 - frequency 401
 - full name 80
 - logging 108
 - login name 80
 - number of retries 26, 512
 - organizations 88
 - password 37, 67, 80
 - phone number 80
 - properties 105
 - provisioning policies 449
 - roles 84
 - SNMP trap listener 18
 - status 460
 - time 123, 139
 - time-out period 26, 512
 - TNS Listener Port 66
 - tools 214
 - topology 425
 - user account 79
 - user name 37
 - user preferences 81
 - user profile 80
 - utilization policies 446
 - zone set 251
- changing ports
 - zone alias 246
- Chargeback 495
 - asset record 459
 - custom information 466
 - infrastructure cost 463
 - licensing information 466
- Chargeback Manager

- about 457
- asset information 463
- asset-based 457, 468
- custom properties 466
- customizing filters 493
- department views 482
- double declining balance 479
- e-mail schedule 486
- e-mailing reports 485
- filtering assets 493
- fixed declining balance 478
- general information 464
- geographic information 466
- hiding filters 494
- reports 484
- saving 461
- schedules editing 488
- selecting element type 492
- setting up 458, 468, 471
- staff information 465
- storage subsystem reports 484
- storage tiers 462
- storage-based 457, 471
- straight line depreciation 477
- viewing by element 481
- viewing by enterprise 483
- viewing report schedules 490
- warranty information 466
- chart
 - bar 401
 - clearing 401
 - filtering data 401
 - line 401
 - printing 401
 - saving 401
 - summary 401
 - time period 401
- charts
 - custom periods 407
 - removing data 406
 - summary 404
 - thresholds 406
 - viewing 404
- checker
 - scan 116
- checking
 - database 151
- child organizations 71
- choosing
 - fabrics 194
- CIM xix
 - log levels 108
- CIM Extension
 - port 503
- CIM Extensions 1
- CIM Object Manager 529
 - tracing 108
- cimom.CimXmlClientHttpConnectTimeout 517
- cimom.ClariionEventPollInterval 330
- cimom.emc.skipRefresh 41
- cimom.hds.exclude 32
- cimom.log 106
- cimom.provider.hds.hiCommandTimeout 285
- cimom.symmetrix.exclude 30
- CIO role 71
- CISCO switches
 - topology 163
 - VSAN 163
- CLARiiON 256, 283
 - polling interval 330
- CLARiiON storage systems
 - no data 517
- cleanup
 - reports 127
- clearing
 - elements 26
 - events 110, 112, 233, 327, 328
 - graph 401
- clearing events 111
- closing
 - Topology windows 187
- CNT
 - switches 17
- cold backup 156
- collectors
 - for reports 236
 - managing 137
 - performance 139
 - reports 121, 123
 - setting time 123
 - starting 114, 121, 340
 - stopping 114, 121
- command window
 - log file 107

- commands
 - custom 204, 207
 - remote 176
- Common Information Model (CIM) 529
- comparing
 - element performance 404
- components 4
- configuring
 - business cost 176
 - Chargeback Manager 458, 468, 471
 - e-mail notification 99, 511
- console
 - buttons 211
 - remote 211
- controller
 - removing 51, 67
- conventions
 - document xx
 - naming zones 316
 - text symbols xx
- copying
 - zone sets 252
- copying text 213
- costs
 - assigning 458
 - business 192
- CR 213
- CRC Errors 408
- creating
 - asset record 459
 - custom commands 204
 - custom name 221
 - custom reports 353
 - departments 467
 - discovery policies 441, 442
 - e-mail schedule 344
 - event policies 444
 - global reports 130
 - host security groups 276
 - LUSE volumes 285
 - new password 80
 - organizations 86
 - roles 83
 - storage pool 261
 - storage volume 265
 - user accounts 78
 - utilization policies 438
 - virtual application 176
 - virtual applications 232
 - volumes 283
 - zone alias members 245
 - zone in fabric 247
 - zone sets 249
- CRLF 213
- cumulative licenses 96
- Current View combo box 425
- custom
 - periods 407
- custom commands
 - about 204
 - adding 204
 - deleting 207
 - editing 207
 - setting up 176
 - stopping 204
- custom name
 - truncated 197
- custom properties 466
- custom reports 350
 - integrating 362
- customer information
 - adding 466
- Customizing 105
- customizing
 - element filters 335
 - filters 334, 335
 - properties 105
 - severity filters 334
- cut off
 - custom name 197
- CXFS 237, 238

D

- data
 - backing up 150
 - performance 400
 - trending 435
- data aging statistics 126
- data collection 114
- Data Discovery Collection
 - e-mail notification 511
- database
 - AIX 515
 - backing up 148, 150

- changing passwords 151
- checking status 151
- cold backup 154, 156
- export 152
- importing
 - Database Admin Utility
 - importing database 153
- reinitializing 153, 158
- restarting 154
- RMAN restore 156
- updating 45
- Database Admin Utility
 - accessing 150
 - archive destination 157
 - archive mode 155
 - changing passwords 151
 - cold backup 156
 - database passwords 151
 - exporting database 152
 - refreshing 150, 157
 - reinitializing database 153
 - resetting logs 158
 - resetting tablespace 154
 - restarting database 154
 - restoring cold backup 154
 - RMAN restore 156
 - temp tablespace 154
 - viewing logs 157
- database connection failed
 - error 514
- date
 - setting 123, 139
- DCOM
 - unable to communicate 514
- deactivating
 - policies 453
 - zone sets 249
- definitions
 - ports 218
- deleting
 - custom commands 207
 - departments 467
 - domain controller 51, 67
 - element 176
 - elements 26, 38, 190
 - e-mail schedules 349, 489
 - events 113, 233, 328, 329
 - fabrics 176, 200
 - host security groups 281
 - license 98
 - multiple jobs 316
 - organizations 88
 - performance data 406
 - policies 453
 - roles 85
 - storage pool 263
 - storage pools 262
 - switches 44
 - TNS Listener Port 66
 - user accounts 80
 - volumes 268, 283
 - zone aliases 246
 - zone sets 249, 251, 523
 - zones 249, 523
- deleting events 111
- department ownership 467
 - removing 475
- departments
 - adding 467
 - deleting 467
 - editing 467
 - removing 467
- designing
 - custom reports 353
- detailed
 - tracing 108
- detailed schema 364
- details
 - event 326
 - obtaining 39
- detecting
 - IP range 15
 - McDATA switches 44
 - switches 44
- device
 - about 530
 - finding 282
 - inaccessible 282
- device issues 524
- devices
 - deleting 38, 190
 - updating 189
- different
 - Java applet 521

- different version
 - Java applet 521
- direct attached storage 418
- directory lookup 408
- disabling
 - discovery schedule 102
 - schedule 102
- discovered
 - hosts 184, 197
 - storage systems 199
- discovered address
 - modifying 37
- Discovered Elements
 - deleting elements 190
- discovered elements
 - deleting elements 38
- discovering
 - applications 51
 - Brocade switches 15, 37
 - CNT switches 17
 - EMC Solutions Enabler 29
 - event details 326
 - HDS storage systems 31
 - HDS systems 32
 - HP XP storage systems 33
 - IP address 15
 - McDATA switches 18
 - Microsoft Exchange 51, 63, 515
 - NetApp filers 34
 - Oracle 51, 52
 - Oracle clusters 52
 - policies 448
 - SQL servers 57
 - storage system 13
 - storage systems 31, 34
 - Sun StorEdge storage systems 35, 36
 - Sun StorEdge switches 17
 - switches 13, 15
 - Sybase 51, 60
 - Symmetrix systems 30
 - troubleshooting 515, 516, 527
- discovery
 - authentication errors 515
 - quarantine 46, 47
 - time-out 517
 - troubleshooting 514
- discovery groups 37
- discovery policies
 - creating 441
- discovery requirements
 - Oracle 515
- discovery schedule
 - adding 100
 - disabling 102
 - editing 102
 - removing 102
- discovery settings
 - importing 45
- discovering
 - IBM storage systems 36
- disk drive 68, 104, 517
- Disk Read 408
- disk space monitoring
 - results of 104
- Disk Utilization 408
- Disk Write 408
- displaying
 - deleted Oracle instances 515
 - element home page 176
 - element impact 190
 - element properties 219
 - filters 335
 - log file 107
 - port details 176
- DNS 515
- document
 - conventions xx
 - prerequisites xix
 - related documentation xix
- documentation, HP web site xix
- Domain Administrator role 71
- domain controller
 - access 66
 - accessing 51, 66
 - removing 51, 67
- domain controller access 51, 66
- double declining
 - calculating 479
- dragging
 - elements 425
- drivers
 - fixing 523
- drives
 - Microsoft Exchange 515

- uninitialized 517
- duplication
 - zone sets 254
 - zones 254

E

- edit
 - event policies 451
- editing
 - cache settings 260
 - custom commands 207
 - date 123
 - departments 467
 - discovery policies 448
 - discovery schedule 102
 - e-mail address 80
 - e-mail schedule 347
 - e-mail schedules 125, 488
 - fabric name 176, 200
 - frequency 401
 - full name 80
 - host security groups 279
 - logging 108
 - login name 80
 - organizations 88, 89
 - ownership percentage 475
 - password 80
 - phone number 80
 - provisioning policies 449
 - roles 84
 - status 460
 - time 123
 - user account 79
 - user preferences 81
 - user profile 80
 - utilization policies 446
 - zone set 251
 - zone sets 249
- editing ports
 - zone alias 246
- EFC Manager 18, 523
- element
 - about 530
 - asset attributes 234
 - deleting 176
 - impact 190
 - performance 400

- refresh 176
- element created in Chargeback
 - about 530
- element details
 - obtaining 39
- element home page
 - viewing 176
- element name
 - filtering 335
- element performance
 - comparing 404
- element properties
 - viewing 219
- element topology 176
- element type
 - filtering 335
 - selecting 335
- elements
 - adding 86, 88
 - arranging 185
 - capacity 428
 - deleting 26, 38, 190
 - dragging 425
 - getting details 45
 - managing 88
 - modifying 37
 - organization 88
 - removing 89
 - topology 185
 - unable to find 516
 - unmanaged 530
 - updating 189
- e-mail
 - notification 454
 - reports 343
 - schedules 125, 343
- e-mail address
 - changing 80
- e-mail notification
 - configuring 99
 - Get Details 511
- e-mail schedule
 - adding 344
 - editing 347
 - organizations 344
 - viewing 490
- e-mail schedules 350

- deleting 349, 489
- editing 488
- viewing 349
- e-mailing
 - chargeback reports 485
 - reports 486
- EMC CLARiiON 31
- EMC Navisphere Agent 283
- EMC Solutions Enabler 29
- EMC Symmetrix 256, 283
- Empty Chart message 433
- error
 - database connection failed 514
 - error -56 522
- Error 503 502
- error message
 - exclusive lock 527
- errors
 - authentication 515
- event
 - details 326
 - icons 194
- Event Manager
 - about 323
 - clearing 327
 - clearing events 111
 - deleting events 111
 - events 327, 329, 331, 334
 - filtering 334
 - SNMP traps 323
 - sorting 329
 - switches 331
- event policies
 - creating 444
- event severity
 - setting 454
- event types
 - all 335
 - all but management server 335
 - applications 335
 - fabrics 335
 - hosts 335
 - management servers 335
 - other 335
 - selecting multiple 335
 - storage systems 335
 - switches 335
- events
 - clearing 110, 112, 233, 327, 328
 - deleting 113, 233, 328, 329
 - excluding 332
 - filtering 333, 334, 335
 - filters 233
 - managing 333
 - policies 451
 - polling 330
 - removing 327
 - severity 233
 - severity levels 333
 - sorting 233, 329
 - switches 331
 - unclearing 233, 327
 - viewing 233
- Events tab
 - about 233
- exceptions 522
- excluding
 - events 332
 - HDS systems 32
 - switches 25
 - Symmetrix systems 30
- exclusive lock
 - error message 527
- export
 - database 152
- external tools 214
 - setting up 214

F

- fabric
 - name 200
 - properties 221
- fabrics
 - changing name 176
 - filtering 194, 425
 - removing 176, 200
 - viewing 221
- file extension
 - assigning 44, 69
- File Server SRM 529
 - properties 115
 - scan checker 116
- File SRM scan 118
- file systems 114

- about 238
- files
 - all logs 106
- filtering
 - assets 493
 - element type 335
 - events 333, 334, 335
 - fabrics 194, 425
 - global reports 342
 - organizations 90
- filtering data 401
- filters
 - Chargeback Manager 493
 - hiding 335, 494
 - host dependency 318
 - LUN dependency 318
 - showing 335
 - storage system dependency 318
 - volume dependency 318
 - zone dependency 319
- Final Destination 408
- finding
 - applications 51
 - capacity 428
 - devices 282
 - element impact 190
 - event details 326
 - hosts 51
 - IP address 15
 - IP range 15
 - performance data 400
 - storage systems 13
 - switches 13
 - trending information 435
- first time
 - users 4
- fixed declining
 - balance 478
- fixing
 - drivers 523
- formatting
 - reports 341
- forwarding
 - SNMP traps 323
- Free Physical Memory 408
- frequency
 - changing 401

- full name
 - changing 80

G

- ganizations 88
- general information
 - adding 464
- generic
 - hosts 184
- generic tape drive 222
- geographic information
 - adding 466
- getting
 - element details 39
- getting details 39, 45
 - applications 51
 - hosts 51
- Global Reporter server
 - about 530
- global reporting view 530
- global reports
 - filtering 342
 - setting up 130
 - tnsnames.ora file 136
- global view 425
 - using 187, 229
- graph
 - bar 401
 - clearing 401
 - filtering data 401
 - line 401
 - printing 401
 - saving 401
 - time period 401
- graphs
 - saving 400
- grey screen 163
- grouping
 - hosts 197
 - storage systems 199
- groups
 - about 197

H

- HBA Risk Analysis 499
- HBAs

- replacing 496
- swapping 523
- HDS 256
 - capacity definitions 428
 - timeout 285
- HDS 9200 285
- HDS 9900 285
- HDS 9900V 285
- HDS storage systems
 - discovering 31
- HdsSkipRefresh 43
- Help button
 - remote console 212
- Help Desk role 71
- help, obtaining xxi
- hiCommandTimeout 285
- hiding
 - filters 335, 494
- hiding tabs 194
- hierarchy
 - organizations 71
- host
 - not in topology 516
- host bus adapter
 - unable to detect 517
- host bus adapters
 - replacing 496
- host dependency
 - host 318
- host persistent binding 530
- host security groups
 - accessing 275
 - creating 276
 - deleting 281
 - editing 279
- hosts
 - adding 302, 311
 - browsing 214
 - discovered 184
 - discovering 51
 - generic 184
 - grouping 197
 - removing 26
 - telnet to 214
 - ungrouping 198
- hot-swapped
 - drives 517

- HP
 - authorized reseller xxi
 - storage web site xxi
 - Subscriber's choice web site xxi
 - technical support xxi
- HP XP storage systems 33
- HTTP Error 503 502
- HTTPS 9
- I
- IBM storage systems
 - discovering 36
- icons
 - event 194
- impact
 - element 190
 - showing 176
- importing
 - database 153
 - discovery settings 45
 - license 96
- inaccessible
 - device 282, 522
- increasing
 - memory 522
- information
 - obtaining element 39
 - performance 400
- infrastructure cost
 - setting 463
- initiator WWN 530
- In-memory Sort Ratio 408
- inode file
 - about 530
- installing
 - Java plug-in 8
 - security certificate 9
- integrating
 - custom reports 362
- internal
 - drives 517
- intervals
 - events polling of 330
- Invalid CRC Errors 408
- IP range 15
- issues
 - devices 524

J

- Java applet
 - different version 521
- Java plug-in 522
 - installing 8
- jboss.properties
 - modifying 105
- jobs
 - provisioning 315

K

- keeping active
 - remote console 211
- killing
 - commands 204
 - custom commands 204

L

- layout
 - arranging 185
- levels
 - event 331
 - severity 333
- LF 213
- Library Cache Hit Ratio 408
- license 93
 - deleting 98
 - importing 96
 - viewing 96, 97
- licensing information
 - adding 466
- line chart 401
- Link Failures 408
- List tab 164, 167
- local drives 515
- locating
 - storage systems 13
 - switches 13
- log file
 - displaying 107
- log messages
 - viewing 26
- logging
 - changing 108
 - viewing 110
- logging off 11

- login name
 - modifying 80
- logs
 - accessing 106
 - refreshing 157
 - resetting 158
 - saving 106
 - viewing 157
- LUN numbering
 - requirements 310
- LUN Security
 - Path Provisioning 299
- LUSE volumes
 - cannot create 285

M

- magnifying
 - Capacity Manager 425
 - topology 425
- MALs 93, 531
- Managed Access Points
 - about 531
- managed access points 93, 531
- managed application license 93, 531
- managed object 531
- management server
 - about 1
 - assigning a name 221
 - backup 150
 - components 4
 - security 71
 - stopping 11
 - variables 207
- managing
 - collectors 121, 139
 - elements 86, 88, 89
 - events 333
 - performance collectors 137
 - switches 44
- mapped 531
- MAPs 93, 531
 - about 531
- materialized view 531
 - about 531
- materialized views
 - previous 397
- maximizing

- screen space 342
- maximum thresholds 406
- McDATA switches 516
 - adding 44
 - discovering 18
- members
 - zone sets 249
- memory
 - increasing 522
- messages
 - data is late 501
 - OutOfMemoryException 522
- messages waiting 408
- meta device 531
- meta volume 531
- metadata client 238
- metadata server 238
- Microsoft Exchange
 - discovering 51, 63, 515
 - drive M 515
- Microsoft Exchange Services 408
- MIME types 44, 69
- minimum
 - thresholds 406
- missing element 531
- mixed mode authentication 57
- modifying
 - cache block size 287
 - cache settings 260, 287
 - Chargeback 458
 - Chargeback Manager 458
 - custom command 207
 - database 45
 - database passwords 151
 - date 123, 139
 - departments 467
 - discovered address 37
 - discovery policies 448
 - domain controller 51, 66
 - elements 37
 - e-mail address 80
 - e-mail schedule 347
 - e-mail schedules 125
 - event policies 451
 - external tools 214
 - fabric name 176, 200
 - filters 493

- frequency 401
- full name 80
- jboss.properties 105
- logging 108
- login name 80
- organizations 88
- ownership percentage 475
- password 37, 67, 80
- phone number 80
- properties 105
- provisioning policies 449
- roles 84
- schedules 124
- SNMP trap listener 18
- status 460
- time 123, 139
- TNS Listener Port 66
- topology 425
- user account 79
- user name 37
- user preferences 81
- user profile 80
- utilization policies 446
- zone members 248
- zone set 251
- zone sets 249

- modifying ports
 - zone alias 246
- monitoring
 - provisioning jobs 315
- monitoring options
 - about 408
- moving
 - elements 185, 425
 - groups 185
 - topology 425
- multipathing 223, 531

N

- name
 - fabric 200
 - filtering 335
 - storage pool 262
 - storage tiers 458
 - zone sets 249
- naming
 - storage tier 462

- naming conventions
 - for zones [316](#)
- naming organizations [71](#)
- navigating elements [215](#)
- navigation details [176](#)
- Navigation screen
 - about [215](#)
 - accessing [219](#)
- Navigation tab
 - duplication [254](#)
- Navisphere [214](#)
- Navisphere Agent [283](#)
- NetApp filers
 - discovering [34](#)
- netcnfg [29](#)
- nethost [29](#)
- Networking [xix](#)
- new password [80](#)
- new window [342](#)
- New Window option [229](#)
- no data
 - CLARiiON storage systems [517](#)
- nonexistant IP addresses [517](#)
- nonexistent Oracle instances [515](#)
- notification
 - e-mail [99](#), [454](#)
- number of retries
 - changing [26](#), [512](#)

O

- obtaining
 - performance data [400](#)
 - security certificate [9](#)
 - utilization reports [431](#)
- old reports [128](#)
- opening
 - new window [229](#), [425](#)
- opening reports
 - new window [342](#)
- Oracle
 - deleted instances [515](#)
 - discovering [51](#), [52](#)
 - discovery requirements [515](#)
- Oracle TNS Listener Port [66](#)
- organizations
 - about [71](#)
 - adding [86](#)
 - deleting [88](#)
 - editing [88](#), [89](#)
 - elements [86](#), [88](#), [89](#)
 - e-mailing reports [344](#)
 - filtering [90](#)
 - properties [83](#)
 - users [83](#)
 - viewing [87](#)
- organizing
 - topology [197](#)
- outdated reports [128](#)
- OutOfMemoryException [522](#)
- overview
 - SAN zoning [240](#), [242](#)
- owners
 - adding [467](#)
 - editing [467](#)
 - removing [467](#)
- ownership
 - percentage [475](#)
 - removing [467](#)

P

- Pan button [425](#)
- parent organizations [71](#)
- Parse CPU [408](#)
- password
 - changing [37](#), [67](#), [79](#), [80](#)
- path information
 - unable to find [516](#)
- Path Provisioning
 - adding hosts [302](#), [311](#)
 - deleting jobs [316](#)
 - executing [314](#)
 - filters [318](#)
 - LUN Security [299](#)
 - monitoring jobs [315](#)
 - storage system dependency filter [318](#)
 - Volume Creation [298](#)
 - Volume Creation and LUN Security [308](#)
 - Volume Creation, LUN Security, and Zone Operation [293](#), [310](#)
 - volume dependency filter [318](#)
 - zone dependency filter [319](#)
 - Zone Operation [304](#)
- Path tab [164](#), [175](#), [423](#)
- paths [223](#)

- percentage
 - ownership 475
- performance
 - collectors 139
 - finding 400
 - statistics 140
- performance collectors
 - managing 137
 - starting 138
 - stopping 139
- Performance Manager 8
 - about 400
 - accessing 400
 - chart format 401
 - charts 400
 - clearing graph 401
 - comparing 404
 - custom periods 407
 - filtering data 401
 - frequency 401
 - multiple elements 401
 - printing 401
 - removing data 406
 - saving graph 401
 - summary charts 404
 - thresholds 406
 - tool bars 401
 - trending 400
- periods
 - custom 407
- persisitent binding 530
- phone number
 - editing 80
- Physical Memory Used 408
- planning organizations 71
- points
 - managed access 93
- policies
 - creating 438
 - deactivating 453
 - deleting 453
 - discovery 441, 448
 - event 444, 451
 - provisioning 442, 449
 - testing 438, 446
 - utilization 438, 446
 - viewing 438, 453
- Policies tab 237
- Policy Manager
 - about 437
 - accessing 438
 - creating policies 437
 - deactivating policies 437, 453
 - deleting policies 437, 453
 - modifying policies 437
 - viewing policies 437, 453
- polling
 - events 330
- pool
 - information 262
 - storage 261
- pools
 - accessing 262
 - deleting 262
- port
 - CIM Extension 503
- port details
 - showing 176
- ports
 - definitions 218
 - viewing 190
- predicting 400
 - performance 400
 - storage capacity 435
- prerequisites xix
- Pre-submission Queue Size 408
- previous materialized views 397
- printing
 - Capacity Manager 431
 - Capacity Manger 425
 - graphs 401
 - System Manager 187
 - topology 230, 425
- privileges
 - roles 71
- problems
 - drivers 523
- process
 - exclusive lock 527
- Processor Utilization 408
- profile
 - user 80
- properties
 - customizing 105

- fabric 221
- File Server SRM 115
- organizations 83
- roles 82
- viewing 219
- provider 531
- provisioning
 - about 239
 - adding hosts 302, 311
 - copying zone sets 252
 - LUN numbering 310
 - monitoring jobs 315
 - policies 449
 - troubleshooting 282, 283, 522, 527
- provisioning jobs
 - scheduling 313
- provisioning policies
 - creating 442
- Provisioning tab 233
- proxy server 531

Q

- quarantine
 - adding elements 46
 - clearing elements 47
- queue size 408

R

- RAID level 262
- Rapid program 37
- read caching 260
- Read Operations 408
- Read Requests 408
- recalculating topology 6
- Receive Queue Size 408
- receiving
 - SNMP traps 323
- refresh
 - element 176
- refresh view 341
- refreshing
 - Database Admin Utility 150
 - element properties 219
 - logs 157
 - report cache 128
 - reports 341

- Symmetrix systems 41
- reinitializing
 - database 153, 158
 - warnings 158
- related documentation xix
- remote
 - commands 176
 - console 211
- remote console
 - buttons 212
 - copying text 213
 - keeping active 211
 - menu options 213
- remote drives 515
- remote scripts
 - running 455
- remote sites
 - filtering 342
- removing
 - custom commands 207
 - department ownership 475
 - departments 467
 - discovery schedule 102
 - domain controller 51, 67
 - element 176
 - elements 26, 38, 88, 89, 190
 - e-mail schedules 349, 489
 - events 113, 233, 327, 328, 329
 - fabrics 176, 200
 - host security groups 281
 - license 98
 - multiple jobs 316
 - organizations 88
 - ownership 467
 - performance data 406
 - policies 453
 - roles 85
 - schedules 124
 - storage pool 263
 - storage pools 262
 - switches 44
 - TNS Listener Port 66
 - user accounts 80
 - volumes 268
 - zone aliases 246
 - zone members 248
 - zone sets 251, 523

- zones 523
- removing ports
 - zone alias 246
- renaming
 - storage tiers 458
- replacing
 - fabric name 200
 - HBA's 496
 - switches 44
- replication 45
- Reporter
 - about 337
 - custom 350
- reports
 - asset based 484
 - Chargeback Manager 486
 - cleanup 127
 - collectors 121
 - custom 350
 - data aging statistics 126
 - deleting e-mail schedules 349
 - designing 353
 - detailed schema 364
 - e-mail 343
 - e-mail schedule 350
 - e-mail schedules 124, 125, 344, 349
 - e-mailing 485
 - formatting 341
 - global 130, 342
 - HTML 341
 - integrating 362
 - Microsoft Excel 341
 - new window 342
 - organizations 344
 - PDF 341
 - refreshing 341
 - refreshing cache 128
 - screen space 342
 - sending 343
 - starting collectors 340
 - storage based 484
 - storage subsystem 484
 - utilization 431
 - XML 341
- Reports tab
 - about 236
- Requests Serviced 408

- resetting
 - tablespace 154
- restarting
 - database 154
- restoring
 - cold backup 154
 - database 151
 - RMAN 156
- right button 428
- risk analysis
 - setting up 499
- RMAN
 - destination 157
- RMAN backup
 - viewing results 149
- RMAN backups 148
- RMAN restore 156
- roles
 - about 71
 - adding 83
 - Application Administrator 71
 - CIO 71
 - deleting 85
 - Domain Administrator 71
 - editing 84
 - Element Control privilege 71
 - Full Control privilege 71
 - Help Desk 71
 - privileges 71
 - properties 82
 - Server Administrator 71
 - Storage Administrator 71
 - users 82
 - View privilege 71
- running
 - cold backup 156
 - remote scripts 455
- Rx Bytes 408

S

- SAN xix, 532
- SAN Zoning Overview 240
- saving
 - chargeback information 461
 - graph 401
 - graphs 400
 - logs 106

- settings to a file 45
 - topology 164, 425
- scan checker 116
- scan status 116
- scanning
 - IP range 15
- schedules
 - deleting 349
 - disabling 102
 - editing 488
 - e-mail 125, 343, 349
 - modifying 124
 - removing 124
- scheduling
 - database backup 148
 - discovery 100, 102
 - File Server SRM scan checker 116
 - provisioning jobs 313
 - report cleanup 127
 - reports 486
- screen space
 - maximizing 342
- scripting
 - variables 207
- scripts
 - remote 455
 - stopping 204
- searching
 - topology 425
- security
 - Management server 71
 - roles 83, 84
- security certificate
 - installing 9
- seeing
 - license 96, 97
- selecting
 - element type 335
 - fabrics 194
 - severity levels 333
- Send Queue Size 408
- sending
 - reports 343
- server
 - backup 150
- Server Administrator role 71
- services
 - stopping 11
- sets
 - zone 249
- setting
 - archive mode 155
 - business cost 176
 - date 123, 139
 - event severity 454
 - infrastructure cost 463
 - report collectors 236
 - storage cost 462
 - storage tier 462
 - thresholds 406
 - time 123, 139
 - time period 401
- setting up
 - Chargeback Manager 458, 468, 471
 - e-mail notification 99
 - external tools 214
 - global reports 130
 - remote commands 176
- settings
 - cache 287
- severity
 - events 233
- Severity combo box 425
- severity levels 194, 331
 - all 333
 - all but clear 333
 - clear 333
 - critical 333
 - informational 333
 - major 333
 - minor 333
 - unknown 333
 - warning 333
- showing
 - element impact 190
 - filters 335
 - impact 176
 - port details 176
- signing out 11
- size
 - cache block 287
 - storage pool 262
- SMTP 99, 532
- SNIA specification 532

- SNMP
 - authentication errors 515
 - trap forwarding 323
- SNMP trap destinations 99
- SNMP trap listener
 - changing 18
- sorting
 - events 233, 329
- specifying
 - custom periods 407
- SQL Server
 - authentication modes 57
- SQL servers
 - discovering 57
- staff information
 - adding 465
- Standard XML Connector 1
- starting
 - collectors 114, 121, 340
 - performance collectors 138
- Statistics 104
 - data aging 126
 - performance 140
- statistics 68
- status
 - changing 460
 - scan 116
- stop button
 - remote console 212
- stopping
 - collectors 114, 121
 - performance collectors 139
 - SAN details 41
 - scripts 204
 - services 11
- storage
 - assigning costs 458
- Storage Administrator role 71
- storage area network 532
- storage based
 - reports 484
- storage capacity
 - predicting 435
- storage cost
 - setting 462
- storage elements
 - viewing 182
- storage pool 532
 - creating 261
 - deleting 263
- storage pools
 - accessing 262
 - deleting 262
- storage subsystem based
 - reports 484
- storage system proxy 531
- storage systems 34, 68, 104
 - discovering 13, 34
 - grouping 199
 - polling interval 330
 - removing 26
 - ungrouping 200
- storage tier
 - setting 462
- storage tiers 458
- storage-based
 - Chargeback Manager 457
- storage-based chargeback
 - setting up 471
- Store Group Size 408
- Store Size 408
- straight line depreciation
 - calculating 477
- Subscriber's choice, HP xxi
- suggested topics 4
- summary text
 - filtering 335
- Sun StorEdge
 - SNMP trap listener 18
- Sun StorEdge A5000 516
- Sun StorEdge storage systems 35, 36
- Sun StorEdge switches 17
- swapped
 - drives 517
- swapping
 - switches 44
- swapping HBAs 523
- switch
 - events 331
- switch port zoning 532
- Switch Risk Analysis 499
- switches
 - adding 44
 - browsing 214

- CISCO 163
 - discovering 13, 15
 - duplication 254
 - excluding 25
 - managing 44
 - McDATA 18, 44, 516
 - name truncated 197
 - number of retries 26, 512
 - port status 218
 - removing 26, 44
 - replacing 44
 - telnet to 214
 - time-out period 26, 512
 - unable to monitor 516
- Sybase
 - discovering 51, 60
- symbols in text xx
- System Event Times 408
- System Manager
 - about 161
 - accessing 163
 - buttons 164
 - can't access 522
 - Cisco switches 163
 - deleting elements 38, 190
 - Discovered hosts 184
 - element impact 190
 - event icons 194
 - global views 187, 229
 - groups 197
 - hiding tabs 194
 - inaccessible 161, 163
 - installing Java plug-in 8
 - Java plug-in 8
 - printing 187
 - scripting variables 207
 - storage elements 182
 - Topology windows 187
 - update 189
 - user interface 164
 - viewing ports 190
 - virtual applications 232
- T
- tablespace
 - resetting 154
- tabs
 - Access 169
 - hiding 194
 - List 167
 - Path 175
- tape drive
 - adding 222
- task
 - running on element 417
- technical support, HP xxi
- telnet to
 - hosts 214
 - switches 214
- testing
 - policies 438, 446
- text
 - filtering 335
- text copying 213
- text symbols xx
- threads 115
- thresholds 406
- tiers
 - storage 458
- time
 - setting 123, 139
- timeout
 - hds 285
- time-out period
 - changing 26
- TNS Listener Port
 - changing 66
- tnsnames.ora file 136
- toolbar
 - Capacity Manager 425
 - System Manager 164
- tools
 - external 214
- topology
 - AIX 515
 - changing 425
 - Cisco switches 163
 - filtering 425
 - host not appearing 516
 - moving 425
 - new window option 425
 - organizing 197
 - printing 230, 425
 - saving 425

- updating 6
- VSANs 163
- Topology tab
 - about 223
- Total Cost/Month for Each Departments 484
- tracing
 - detailed 108
- trap forwarding 323
- trap generation 99
- traps
 - SNMP 323
- trending
 - Capacity Manager 435
 - reports 337
- triangle
 - icons 194
- troubleshooting
 - CLARiiON storage systems 283
 - discovery 514
 - discovery and getting element details 514, 515, 516, 527
 - EMC storage systems 283
 - EMC Symmetrix storage systems 283
 - Engenio storage systems 287
 - HDS storage systems 285
 - JavaScript blockers 163
 - Microsoft Exchange 515
 - popups 163
 - provisioning 282, 283, 522, 527
- truncated
 - custom name 197
- Tx Bytes 408
- types of
 - monitoring 408

U

- unable to
 - discover 514
- Unable to access resource 522
- unable to detect
 - host bus adapter 517
- unable to find
 - elements 516
 - path information 516
- unable to retrieve data 523
- unclearing
 - events 233, 327

- ungrouping
 - hosts 198
 - storage systems 200
- uninitialized
 - drives 517
- updating
 - database 45
 - element properties 219
 - elements 189
 - license 96
- updating topology 6
- used pools 262
- user accounts
 - creating 78
 - deleting 80
- user audit log 107
- user interface
 - System Manager 164
 - about 164
- user name
 - changing 37
- user preferences
 - changing 81
- user profile
 - modifying 80
- users
 - about 71
 - adding 78
 - first time 4
 - organizations 83
 - roles 82, 83, 84
- using
 - global view 187, 229
- utilization
 - policies 446
- utilization policies
 - creating 438
- Utilization tab 423

V

- variables
 - scripting 207
- viewing
 - all logs 106
 - asset attributes 461
 - capacity 428
 - capacity charts 433

- chargeback by element 481
- charts 404
- cumulative licenses 96
- data aging statistics 140
- database status 151
- department views 482
- element asset attributes 234
- element home page 176
- element impact 190
- element properties 219
- e-mail schedules 349, 490
- event details 326
- events 233
- fabrics 194, 221
- log file 107
- log messages 26
- logging 110
- logs 157
- organization properties 83
- organizations 87
- performance data 400
- policies 438, 453
- ports 190
- report refresh 341
- RMAN backup results 149
- schedule history 491
- scheduled e-mail deliveries 124
- security certificate 9
- specific license 97
- storage elements 182
- trending information 435
- utilization reports 431
- volume information 263
- zone aliases 244
- views
 - global 187, 229
- virtual application
 - about 532
 - adding 176, 232
 - creating 232
- Virtual Memory Used 408
- volume
 - meta 531
- volumes
 - creating 283
 - deleting 268, 283
 - unable to create 285

- viewing information 263
- VSAN 163

W

- warnings
 - reinitializing 158
- warranty information
 - adding 466
- WBEM 532
 - about 532
- web sites
 - HP documentation xix
 - HP storage xxi
 - HP Subscriber's choice xxi
- window
 - opening reports 342
- Windows Management Instrumentation 532
- WMI 532
 - about 533
- Worldwide Name 533
- WorldWide Name zoning 532
- Worldwide Name zoning 533
 - about 533
- write caching 260
- Write Operations 408
- WWN 533

X

- XFS 237

Z

- zone
 - hard 533
 - soft 533
- zone alias
 - about 533
 - creating 245
 - deleting 246
 - modifying ports 246
- zone hard
 - about 533
- zone member 533
- zone members
 - adding 248
 - removing 248
- zone set

- creating [249](#)
- zone sets
 - accessing [249](#)
 - activating [253](#)
 - copying [252](#)
 - creating [249](#)
 - deleting [251](#), [523](#)
 - duplication [254](#)
 - editing [249](#)
 - members [249](#)
 - modifying [251](#)
 - removing [249](#)
- zones
 - creating [247](#)
 - deleting [249](#), [523](#)
 - duplication [254](#)
 - naming conventions [316](#)
 - removing [249](#)
 - viewing aliases [244](#)
- Zoning Library
 - duplication [254](#)

Figures

1	Accessing Features from the Tools > Storage Essentials Menu	2
2	Accessing Roles in Storage Essentials	2
3	Accessing the Documentation Center	3
4	Status Light	7
5	Parent-Child Hierarchy for Organizations	75
6	Children in Multiple Organizations	76
7	Changing Your User Profile	80
8	Accessing the User Preferences Tab	81
9	Viewing Organizations	87
10	Clicking the Organization Link	91
11	Active Organization	91
12	An Example of Direct Attached Storage	95
13	Selecting the Date and Time from the Calendar	116
14	About the File System Viewer Scan Checker	117
15	Report Views and Report Cache Refresh	119
16	Selecting the Date and Time from the Calendar	124
17	An Example of Global Reporting	131
18	Selecting the Date and Time from the Calendar	140
19	Expanding the Fabric Node	168
20	Fabric Members Are Highlighted	168
21	Highlighting the Applications in the Topology	169
22	Members of a Zone Set	170
23	Displaying Zone Members	171
24	Zone Member to Switch	171
25	Obtaining Information About a Zone Member's Adapter	172
26	Highlighting Elements with Host Bindings	172
27	Displaying Host Bindings	173
28	HBA Port Properties	173
29	WWN Properties	174
30	WWN Properties	174
31	Obtaining Path Information	175
32	Path Information Visible in the Tree	176
33	Enclosing the Elements	186
34	Dragging Multiple Elements to Their New Location	186
35	Showing the Impact of an Application	191
36	Determining Business Cost	193
37	Expanding the Topology Pane	194
38	Filtering Fabrics	194
39	Remote Console	211
40	Obtaining Information About a Host	215
41	Details of a Host Connected to a Switch	216
42	Finding the Status of a Port	218
43	Topology of a Server	224
44	Topology of a Server (Continued)	225
45	Multipathing Displayed in the Topology	226
46	Multipathing Displayed in the Topology (Continued)	227

47	Direct Attached Storage in the Topology	228
48	New Window Option	230
49	Viewing Asset Records	235
50	CXFS File System	238
51	Resources in Two Zones	241
52	Overview of Zoning Capabilities	242
53	Selecting a Storage System	294
54	Selecting a Storage System	301
55	Selecting a Storage System	305
56	Selecting a Storage System	309
57	Specifying a LUN Number	311
58	Deleting Multiple Provisioning Jobs	316
59	Hiding the Middle Pane	342
60	Selecting Organizations Used in This Report	346
61	Selecting Organizations Used in This Report	348
62	Report Architecture	351
63	Choosing a Standard Report	354
64	Adding Tables for a Standard Report	355
65	Linking Common Data in Tables for a Standard Report	356
66	Creating Search Criteria for Standard Reports	357
67	Deciding Which Data Should Appear in the Report	358
68	Sorting Information in the Report	359
69	Selecting the Layout of the Report	360
70	Report Template Displayed	361
71	Clicking the View Tab	362
72	Legend for Capacity Manager	424
73	Capacity of an Element	425
74	Viewing the Capacity of Elements in a Fabric	428
75	Viewing a Utilization Report	431
76	Chargeback Capacity	476
77	Capacity in Capacity Manager	476

Tables

1	Document conventions	xx
2	Status Light Settings	7
3	Discovery Requirements for Switches	14
4	Required Switch Models and InVsn Versions for Discovery16	
5	Discovery Settings for McDATA and Connectrix Switches	19
6	Discovery Requirements for Storage Systems and NAS Filers	27
7	Default Role Privileges	72
8	Default Role Privileges with Elements	73
9	Changing User Preferences for Event Monitoring for Storage Essentials	81
10	License Requirements	93
11	Determining Managed Access Points	94
12	Logging Levels	108
13	Default Settings for Clearing Events	112
14	About the Collectors for File System Viewer	114
15	About Collectors for Reports	121
16	Viewing System Wide E-mail Schedules	124
17	Description of "An Example of Global Reporting"	131
18	About Performance Collectors	138
19	Description of Backups	144
20	Backup Directories in %ORA_HOME%\rman146	
21	Backup Example	147
22	Default Passwords for Database Accounts	158
23	Description of Toolbar in System Manager	164
24	Elements Displayed in the Topology	167
25	Drop-Down Menu Options Accessible from the Topology*	177
26	Drop-Down Menu Options on the Access Tab	181
27	Show Impact Results	192
28	Legend for Severity Level Icons	195
29	Descriptions of the Severity Levels Presented in the Severity Menu195	
30	Variables for All Elements*	208
31	Variables for Only Storage Systems, Switches, and Hosts*	208
32	Variables for Only Switches	209
33	Variables for Only Hosts*	209
34	Variables for Only Applications	210
35	Buttons on the Remote Console	212
36	Drop-Down Menu Options	213
37	Information Available from the Navigation Page	216
38	Port Status Definitions	218
39	About the Toolbar in the Topology Tab	229
40	Setting Up Zoning	242
41	Zoning Support	243
42	Provisioning and Pool Support	257
43	Volume and HSG Support	258

44	Setting Up Storage Partitioning	259
45	HSG Port Assignments	270
46	Allowed Initiators in Host Security Groups	272
47	Volume Usage.	288
48	Overview for Path Provisioning	292
49	Zone Icons	297
50	Zone Icons	307
51	Zone Icons	321
52	About the Icons in Event Monitoring for Storage Essentials	324
53	Supported Hardware	324
54	Default Settings for Clearing Events	328
55	Brocade Switch Events	331
56	Severity Levels.	333
57	Viewing E-mail Schedules for a Report	349
58	Description of the Report Views	364
59	MVC_HOSTSUMMARYVW	367
60	MVC_CARDSUMMARYVW	368
61	MVC_HOSTVOLUMESUMMARYVW (logical volumes).	369
62	MVC_HOSTDISKDRIVEVW	369
63	MVC_STORAGESYSTEMSUMMARYVW	370
64	MVC_STORGAEPOOLSUMMARYVW	370
65	MVC_STORAGEVOLUMESUMMARYVW	371
66	MVC_SWITCHSUMMARYVW	372
67	MVC_PORTSUMMARYVW	374
68	MVC_ZONESUMMARYVW	374
69	MVC_ZONEVW	375
70	MVC_PATHVW.	376
71	MVC_SUBPATHVW.	376
72	MVC_MULTIPATHVW	377
73	MVC_EVENTSVW.	377
74	MVC_ORGANIZATIONVW	378
75	MVC_ORGRELATIONVW.	378
76	MVC_HOSTCAPACITYVW	379
77	MVC_STORAGESYSTEMCONFIGVW	379
78	MVC_STORAGEPOOLCONFIGVW	380
79	MVC_SWITCHCONFIGVW	380
80	MVC_OPTIONALTABLEVW	380
81	MVC_DISKEXTENTSUMMARYVW	381
82	MVC_STORAGEVOLUMEPORTS	381
83	MVC_VOLUMEDISKDRIVEVW	382
84	MVC_STORAGEPROCESSORSUMMARYVW	382
85	MVC_DISKDRIVESUMMARYVW	383
86	MVC_DISK_EXTENTVW	384
87	MVC_ASSETSUMMARY	384
88	MVC_APPLICATIONSUMMARYVW	386
89	MVC_UNITACCESSVW	387
90	MVCA_DBAPPCAPACITYVW	387

91	MVCA_EXCHAPPCAPACITYVW	388
92	MVCA_VIRTUALAPPCAPACITYVW	388
93	MVCA_FSRM_VOLUMESUMMARYVW	389
94	MVCA_FSRM_AGESUMMARYVW	389
95	MVCA_FSRM_EXTDETAILSUMMARYVW	390
96	MVCA_FSRM_DIRDETAILSUMMARYVW	390
97	MVCA_FSRM_USERSUMMARYVW	390
98	MVCA_FSRM_GROUPSUMMARYVW	391
99	MVCA_FSRM_TOPNFILES	391
100	MVCA_FSRM_AGEDFILEDETAILS	392
101	MVCA_FSRM_LARGEDIRINFO	392
102	MVCA_BU_MASTERSERVERSUMMARY	392
103	MVCA_BU_MEDIASERVERSUMMARY	393
104	MVCA_BU_CLIENTSUMMARY	393
105	MVCA_BU_MEDIASUMMARY	394
106	MVCA_BU_JOBSUMMARY	395
107	MVCA_BU_LIBRARYSUMMARY	396
108	MVIEWCORE_STATUS	397
109	MVIEW_STATUS	397
110	Views from Previous Releases	398
111	Toolbar in Lower Pane of Performance Manager	401
112	About the Monitoring Options	408
113	Host Monitoring Support	419
114	Toolbar in Lower Pane of Capacity Manager	425
115	Explanation of the Properties of the Capacity Levels for HDS Array Groups	430
116	Policy Templates	439
117	Policy Table Description	453
118	Setting Up Chargeback Manager	459
119	Element Type Icons	462
120	Viewing E-mail Schedules for a Chargeback Report	490
121	Element Types	492
122	Troubleshooting Firewalls	505
123	Time-out Properties	513
124	Retry Properties	513
125	Known Device Issues	524

